JPL D-11399, Rev. D

Multi-angle Imaging SpectroRadiometer (MISR)

Level 2 Cloud Detection and Classification Algorithm Theoretical Basis

Approval:

David J. Diner
MISR Principal Investigator

The MISR web site should be consulted to determine the latest released version of this document (http://www-misr.jpl.nasa.gov).
Approval signatures are on file with the MISR Project.

Jet Propulsion Laboratory
California Institute of Technology

TABLE OF CONTENTS

GLOSSARY OF ACRONYMS

A

AGP (Ancillary Geographic Product)
ARP (Ancillary Radiometric Product)
ASCM (Angular Signature Cloud Mask)
ATB (Algorithm Theoretical Basis)
ATSR (Along-Track Scanning Radiometer)
AU (Astronomical Unit)
AVHRR (Advanced Very High Resolution Radiometer)
AVIRIS (Airborne Visible and Infrared Imaging Spectrometer)

B

BDAS (Band-Differenced Angular Signature)
BRF (Bidirectional Reflectance Factor)

C

CCD (Charge-Coupled Device)
CSM (Cloud Shadow Mask)
CSSC (Cloud Screening Surface Classification)

D

DAO (Data Assimilation Office)
DSVI (D Spatial Variability Index)

E

ECR (Earth Centered Reference)
ECS (EOSDIS Core System)
EOSDIS (Earth Observing System Data and Information System)

G

GDQI (Geometric Data Quality Indicator)
GLDV (Gray Level Difference Vector)

I

IFOV (Instantaneous Field Of View)

M

M2 (Multipoint Matcher)
M3 (Multipoint Matcher using Medians)

MISR (Multi-angle Imaging SpectroRadiometer)
MODIS (Moderate Resolution Imaging Spectroradiometer)

N

NM (Nested Maxima)
NR (No Retrieval)
NSIDC (National Snow and Ice Data Center)

P

PGS (Product Generation System)

R

RCCM (Radiometric Camera-by-camera Cloud Mask)
RDQI (Radiometric Data Quality Indicator)
RLRA (Reflecting Level Reference Altitude)
RS (Rank Sort)

S

SCF (Science Computing Facility)
SDCM (Stereoscopically Derived Cloud Mask)
SDP (Science Data Production)
SOM (Space Oblique Mercator)

T

TASC (Terrestrial Atmosphere and Surface Climatology)
TOA (Top-of-Atmosphere)
TSM (Topographic Shadow Mask)

U

UTC (Coordinated Universal Time)

1. INTRODUCTION

1.1 PURPOSE

This Algorithm Theoretical Basis (ATB) document describes the algorithms used to retrieve the cloud classification parameters of the MISR Level 2 Top-of-Atmosphere (TOA)/Cloud Product. These parameters are summarized in Table 1. In particular, this document identifies sources of input data, both MISR and non-MISR, required for parameter retrievals; provides the physical theory and mathematical background underlying the derivation of the cloud classification parameters; includes implementation details; and describes assumptions and limitations of the adopted approach. It is used by the MISR Science Data System Team to establish requirements and functionality of the data processing software.

Table 1: Cloud classification parameters in the Level 2 TOA/Cloud Product

Parameter name	Units	Horizontal Sampling and Coverage	Comments
Reflecting-Level Reference Altitude (RLRA)	km	2.2 km (Global)	• Retrieved using MISR stereo imagery
Stereoscopically Derived Cloud Mask (SDCM)	none	1.1 km (Global)	• Retrieved using MISR stereo imagery and Radiometric Camera-by-camera Cloud Mask (RCCM) from stereo reference camera
Angular Signature Cloud Mask (ASCM)	none	1.1 km (Global)	• Retrieved using the Band-Differenced Angular Signature (BDAS) method
TOA bidirectional reflectance factor (BRF)	none	2.2 km (Global)	• Radiance ratioed to that from ideal lambertian reflector at RLRA • 9 viewing angles, 4 spectral bands • Computed independently for radiation leaving tops and sides of RLRA columns • Includes number of unobscured pixels
Texture indices	none	2.2 km (Global)	• From 672-nm statistics, 9 angles
Regional, altitude-binned scene classifiers	none	17.6 km (Global)	• Uses stereo heights to provide altitude binning • Referenced to surface terrain • Low-, middle-, and high-altitudes
Regional scene classifiers	none	17.6 km (Global)	• Derived from all cloud masks
Regional, angle-by-angle cloud fractions	none	17.6 km (Global)	• Derived from the RCCM • 9 angles
Cloud Shadow Mask	none	1.1 km (Global)	• Uses stereo heights, RCCM, and SDCM
Topographic Shadow Mask	none	1.1 km (Global)	• Uses AGP terrain elevations

Table 1: Cloud classification parameters in the Level 2 TOA/Cloud Product (continued)

Parameter name	Units	Horizontal Sampling and Coverage	Comments
Quality assessment indicators	vary	vary	• Include confidence indicators for cloud masks, RLRA standard deviation, and stereo matcher statistics

1.2 SCOPE

This document covers the algorithm theoretical basis for the cloud classification parameters of the TOA/Cloud Product which are to be routinely retrieved at the DAAC. Post-launch and specialized products or parameters are not discussed. Current development and prototyping efforts may result in modifications to parts of certain algorithms. Only the algorithms which will be implemented at the DAAC for routine processing will be preserved in the final release of this document.

Chapter 1 describes the purpose and scope of the document. Chapter 2 provides a brief overview. The processing concept and algorithm description are presented in Chapter 3. Chapter 4 summarizes assumptions and limitations. References for publications cited in the text are given in Chapter 5. Literature references are indicated by a number in italicized square brackets, e.g., *[1]*.

1.3 MISR DOCUMENTS

Reference to MISR Project Documents is indicated by a number in italicized square brackets as follows, e.g., *[M-1]*. The MISR web site (http://www-misr.jpl.nasa.gov) should be consulted to determine the latest released version of each of these documents.

[M-1] Experiment Overview, JPL D-13407.

[M-2] Data Product Description, JPL D-11103.

[M-3] Level 1 Radiance Scaling and Conditioning Algorithm Theoretical Basis, JPL D-11507.

[M-4] Level 1 Georectification and Registration Algorithm Theoretical Basis, JPL D-11532.

[M-5] Level 1 Cloud Detection Algorithm Theoretical Basis, JPL D-13397.

[M-6] Level 1 In-flight Radiometric Calibration and Characterization Algorithm Theoretical Basis, JPL D-13398.

[M-7] Level 1 Ancillary Geographic Product Algorithm Theoretical Basis, JPL D-13400.

[M-8] Level 1 In-flight Geometric Calibration Algorithm Theoretical Basis, JPL D-13399.

[M-9] Level 2 Top-of-Atmosphere Albedo Algorithm Theoretical Basis, JPL D-13401.

[M-10] Level 2 Aerosol Retrieval Algorithm Theoretical Basis, JPL D-11400.

[M-11] Level 2 Surface Retrieval Algorithm Theoretical Basis, JPL D-11401.

[M-12] Level 2 Ancillary Products and Datasets Algorithm Theoretical Basis, JPL D-13402.

[M-13] Algorithm Development Plan, JPL D-11220.

[M-14] In-flight Radiometric Calibration and Characterization Plan, JPL D-13315.

[M-15] In-flight Geometric Calibration Plan, JPL D-13228.

[M-16] Science Data Validation Plan, JPL D-12626.

[M-17] Science Data Processing Sizing Estimates, JPL D-12569.

[M-18] Science Data Quality Indicators, JPL D-13496.

1.4 REVISIONS

The original version of this document was dated March 3, 1994. Revision A was released December 19, 1994. Revision B was released August 15, 1996. Revision C was released December 9, 1997. This version is Revision D. Changes from Rev. C are indicated by change bars, as shown at the left.

2. EXPERIMENT OVERVIEW

2.1 OBJECTIVES OF MISR CLOUD CLASSIFICATION

As a result of their large areal extent, high albedo, and variability on many timescales, clouds play a major role in governing the Earth's energy balance. Regional studies of the impact of clouds on the energy balance require measurements of the radiation budgets as a function of scene type. The importance of cloud characteristics in global studies of climate has been well documented *[6]*. Current theories and models of the response of the Earth's climate system to, for example, the increase in trace gases, are severely limited by our present ignorance of the feedback processes associated with changes in cloud amount and cloud properties. In this respect, two issues are paramount. One is the realistic modeling of cloud-radiation interaction taking into account the variable structure of broken cloud fields and processes that occur at the sub-grid scale level of present general circulation models. The other is the ability to invert satellite measured radiances to obtain hemispherical fluxes with sufficient resolution to discriminate between cloud-filled and cloud-free scenes.

Deriving from its ability to measure any scene from multiple directions, MISR will contribute unique cloud classification information. This will enable study, on a global basis, of the effects of different types of cloud fields (classified by their heterogeneity and altitude) on the spectral solar radiance and irradiance reflected to space, including spatial and temporal dependences. The most important elements of the MISR retrievals are accurate spectral albedos and spectral bidirectional reflectance factors, coupled to useful scene information, such as parameterizations of the cloud morphology.

A scientific background and historical perspective on related cloud studies using remote sensing, the unique contributions of MISR, and a scientific rationale for the cloud classification parameter contents of the MISR TOA/Cloud Product are presented in *[M-1]*.

2.2 INSTRUMENT CHARACTERISTICS

The MISR instrument consists of nine pushbroom cameras. It is capable of global coverage every nine days, and flies in a 705-km descending polar orbit. The cameras are arranged with one camera pointing toward the nadir (designated An), one bank of four cameras pointing in the forward direction (designated Af, Bf, Cf, and Df in order of increasing off-nadir angle), and one bank of four cameras pointing in the aftward direction (using the same convention but designated Aa, Ba, Ca, and Da). Images are acquired with nominal view zenith angles, relative to the surface reference ellipsoid, of 0°, 26.1°, 45.6°, 60.0°, and 70.5° for An, Af/Aa, Bf/Ba, Cf/Ca, and Df/Da, respectively. Each camera uses four Charge-Coupled Device (CCD) line arrays in a single focal plane. The line arrays consist of 1504 photoactive pixels plus 16 light-shielded pixels per array, each 21 μm by 18 μm. Each line array is filtered to provide one of four MISR spectral bands. The

spectral band shapes are approximately gaussian and centered at 446, 558, 672, and 866 nm.

MISR contains 36 parallel signal chains corresponding to the four spectral bands in each of the nine cameras. The zonal overlap swath width of the MISR imaging data (that is, the swath seen in common by all nine cameras along a line of constant latitude) is ≥ 360 km, which provides global multi-angle coverage of the entire Earth in 9 days at the equator, and 2 days near the poles. The cross-track IFOV and sample spacing of each pixel is 275 m for all of the off-nadir cameras, and 250 m for the nadir camera. Along-track IFOV's depend on view zenith angle, ranging from 214 m in the nadir to 707 m at the most oblique angle. Sample spacing in the downtrack direction is 275 m in all cameras. The instrument is capable of buffering the data to provide 4 sample x 4 line, 2 sample x 2 line, or 1 sample x 4 line averages, in addition to the mode in which pixels are sent with no averaging. The averaging capability is individually selectable within each of the 36 channels, and there are several observational modes of the MISR instrument. The MISR TOA/Cloud Product is generated from Global Mode data. Global Mode refers to continuous operation with no limitation on swath length. Global coverage in a particular spectral band of one camera is provided by operating the corresponding signal chain continuously in a selected resolution mode. Any choice of averaging modes among the nine cameras that is consistent with the instrument power and data rate allocation is suitable for Global Mode. Current plans are to operate the instrument in the 4 x 4 averaging mode (1.1-km sampling) with selected channels operated in 1 x 1 or 1 x 4 mode.

Most of the highest resolution observations will be acquired in the red (672-nm) band, as this is expected to be the wavelength where the imagery will have the highest contrast, based upon considerations of atmospheric haze, land and ocean reflectivity, and instrument performance. These observations are central to the stereoscopic and texture-based approaches to be used as part of MISR cloud classification.

Additional background on the instrument design is provided in *[M-1]*.

2.3 MISR CLOUD CLASSIFICATION STRATEGY

Cloud classification by MISR should help to partition clouds into categories distinguished by parameters such as: (a) cloud elevation, (b) angular signature, and (c) texture or degree-of-brokenness (e.g., stratiform vs. cumuliform).

In order to establish a parameter that serves the purpose of providing a classification scheme that incorporates cloud altitudes, and which also serves as a dynamic (i.e., spatially varying) reference altitude for co-registering the multi-angle views, a reference level known as the Reflecting Level Reference Altitude is established. The RLRA is defined to be the level found by matching features (or areas) with the greatest contrast in the near-nadir viewing directions. Physically, this corresponds to the main reflecting layer, which will typically be either the tops of bright clouds, or under atmospheric conditions corresponding to clear skies or thin cloud, it will be located at the

surface. The RLRA is defined over subregions measuring 2.2 km x 2.2 km. The algorithm for retrieving RLRA is stereophotogrammetric in nature, and makes use of the multiple views from satellite altitude over a wide angular range to separate the effects of wind displacement independently from height. For scene classification purposes, the stereoscopic heights used in deriving the RLRA are also used to establish a Stereoscopically-Derived Cloud Mask (SDCM), and are used to generate regional altitude-binned cloud fraction metrics. The RLRA is also used as a common reference altitude for projecting the multi-angle bidirectional reflectance factors, from which texture indices will be derived. Unique to MISR is the ability to retrieve the angular signature of texture. Several straightforward textural parameters, such as standard deviation divided by the mean and difference-vector standard deviation, will be retrieved as part of the MISR TOA/Cloud Product. The correspondence between these parameterizations and physical cloud type, particularly with respect to distinguishing stratiform and cumuliform clouds, will be the subject of research utilizing these parameters in the post-launch era.

Detection of high cirrus is also important to MISR. The lack of MISR channels beyond 1.0 μm implies that the detection of optically thin cirrus on a per-camera basis will be difficult. Nadir imagers cannot always detect cirrus clouds due to restricted phase angle coverage. Therefore, the multi-angle strategy will be used in a novel way. Motivated by the future availability of MISR observations and the difficulty many current algorithms have in detecting thin cirrus, Di Girolamo *[2]* and Di Girolamo and Davies *[3]* have developed a method known as Band-Differenced Angular Signature (BDAS), a technique which takes advantage of the difference in the Rayleigh signal above high clouds between blue and red or near-IR wavelengths as a function of angle. High clouds have a unique signature that distinguishes them from clear sky and low-level clouds. Application of this technique is used to generate an additional cloud classifier, known as the Angular Signature Cloud Mask (ASCM).

3. ALGORITHM DESCRIPTION

3.1 PROCESSING OUTLINE

Routine in-flight standard processing at the DAAC to carry out the MISR cloud classification occurs in three stages, described below.

3.1.1 Stage 1 (preliminary processing)

Stage 1 consists of these steps:

(1a) Averaging of radiances to the spatial resolutions required for application of the subsequent algorithms.

(1b) Conversion of input radiances from Level 1B2 to bidirectional reflectance factors (BRF's), providing intermediate datasets at the same resolution as the input radiances. One dataset is calculated from terrain-projected TOA radiances and is referenced to the surface topography; the other is calculated from ellipsoid-projected TOA radiances and is referenced to the WGS84 surface ellipsoid.

(1c) Calculation of ellipsoid-referenced look vectors for each camera, on 2.2-km centers, and solar illumination vectors, on 17.6-km centers.

3.1.2 Stage 2 (stereophotogrammetric processing)

Stage 2 consists of these steps:

(2a) Establish which sets of cameras will be used for stereoscopic wind and height retrievals.

(2b) Determination of a stereoscopically-derived velocity and height field using area-based and feature-based stereo matching techniques.

(2c) Establishment of altitude designations associated with the retrieved heights (i.e., near the surface or above the surface), along with a high or low confidence flag.

(2d) Combination of the stereoscopically-derived altitude designations with the Radiometric Camera-by-camera Cloud Mask (RCCM) generated during Level 1B2 processing, according to a specified decision matrix, to generate a height field incorporating all available information.

(2e) Use of the final height field to establish the Stereoscopically-Derived Cloud Mask (SDCM) with 1.1-km sampling, and values of RLRA and standard deviation of RLRA within $(2.2\text{-km})^2$ subregions. The RLRA and SDCM are used to co-register multi-angle views, and to aid in cloud classification.

(2f) Conceptualization of the RLRA surface as an array of square $(2.2\text{-km})^2$ columns (prisms), and projection of BRF's to the tops and sides of these prisms.

(2g) Calculation of the reflecting-level parameters: average BRF's, texture indices, and number of unobscured pixels on the tops of the RLRA prisms, and average BRF's and number of unobscured pixels on the sides of the RLRA prisms.

3.1.3 Stage 3 (scene classification)

Stage 3 consists of these steps:

(3a) Construction of an Angular Signature Cloud Mask (ASCM) using the BDAS algorithm on the B, C and D cameras which view forward-scattered light.

(3b) Updating of the seasonal cirrus detection thresholds based upon histogram analysis.

(3c) Cloud classification by height, using the stereoscopically retrieved heights, as well as the SDCM, ASCM, and information derived from the RCCM.

(3d) Generation of other regional scene classifiers (cloudy and clear-sky fractions).

(3e) Generation of a Cloud Shadow Mask and a Topographic Shadow Mask.

Processing flow concepts are shown diagrammatically throughout the document. The convention for the various elements displayed in these diagrams is shown in Figure 1.

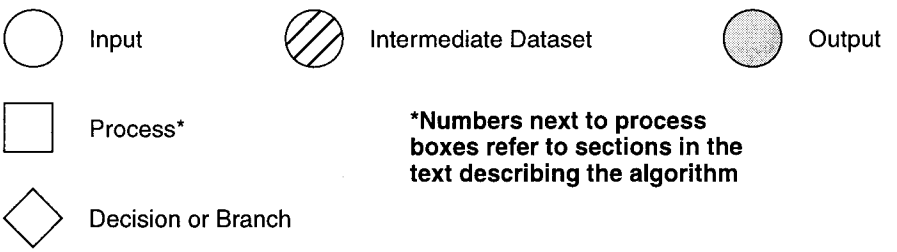

Figure 1. Conventions used in processing flow diagrams

Overviews of the TOA/cloud processing flow concept for each of the three stages is shown in Figures 2 - 7.

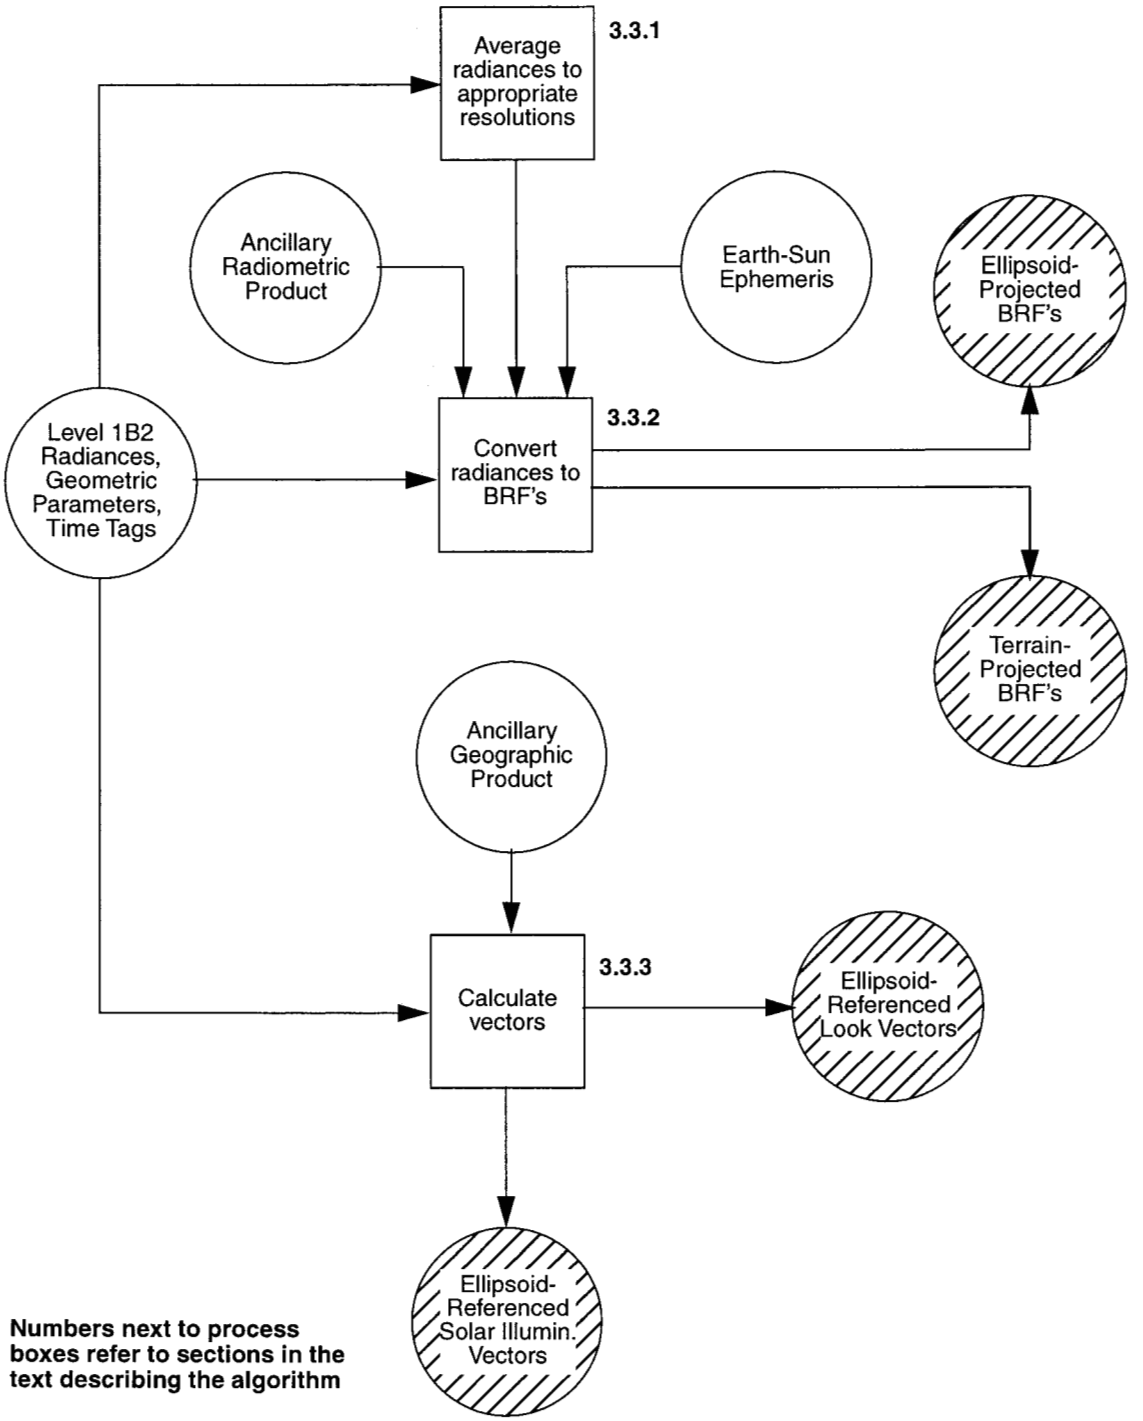

Figure 2. Stage 1 TOA/Cloud processing overview

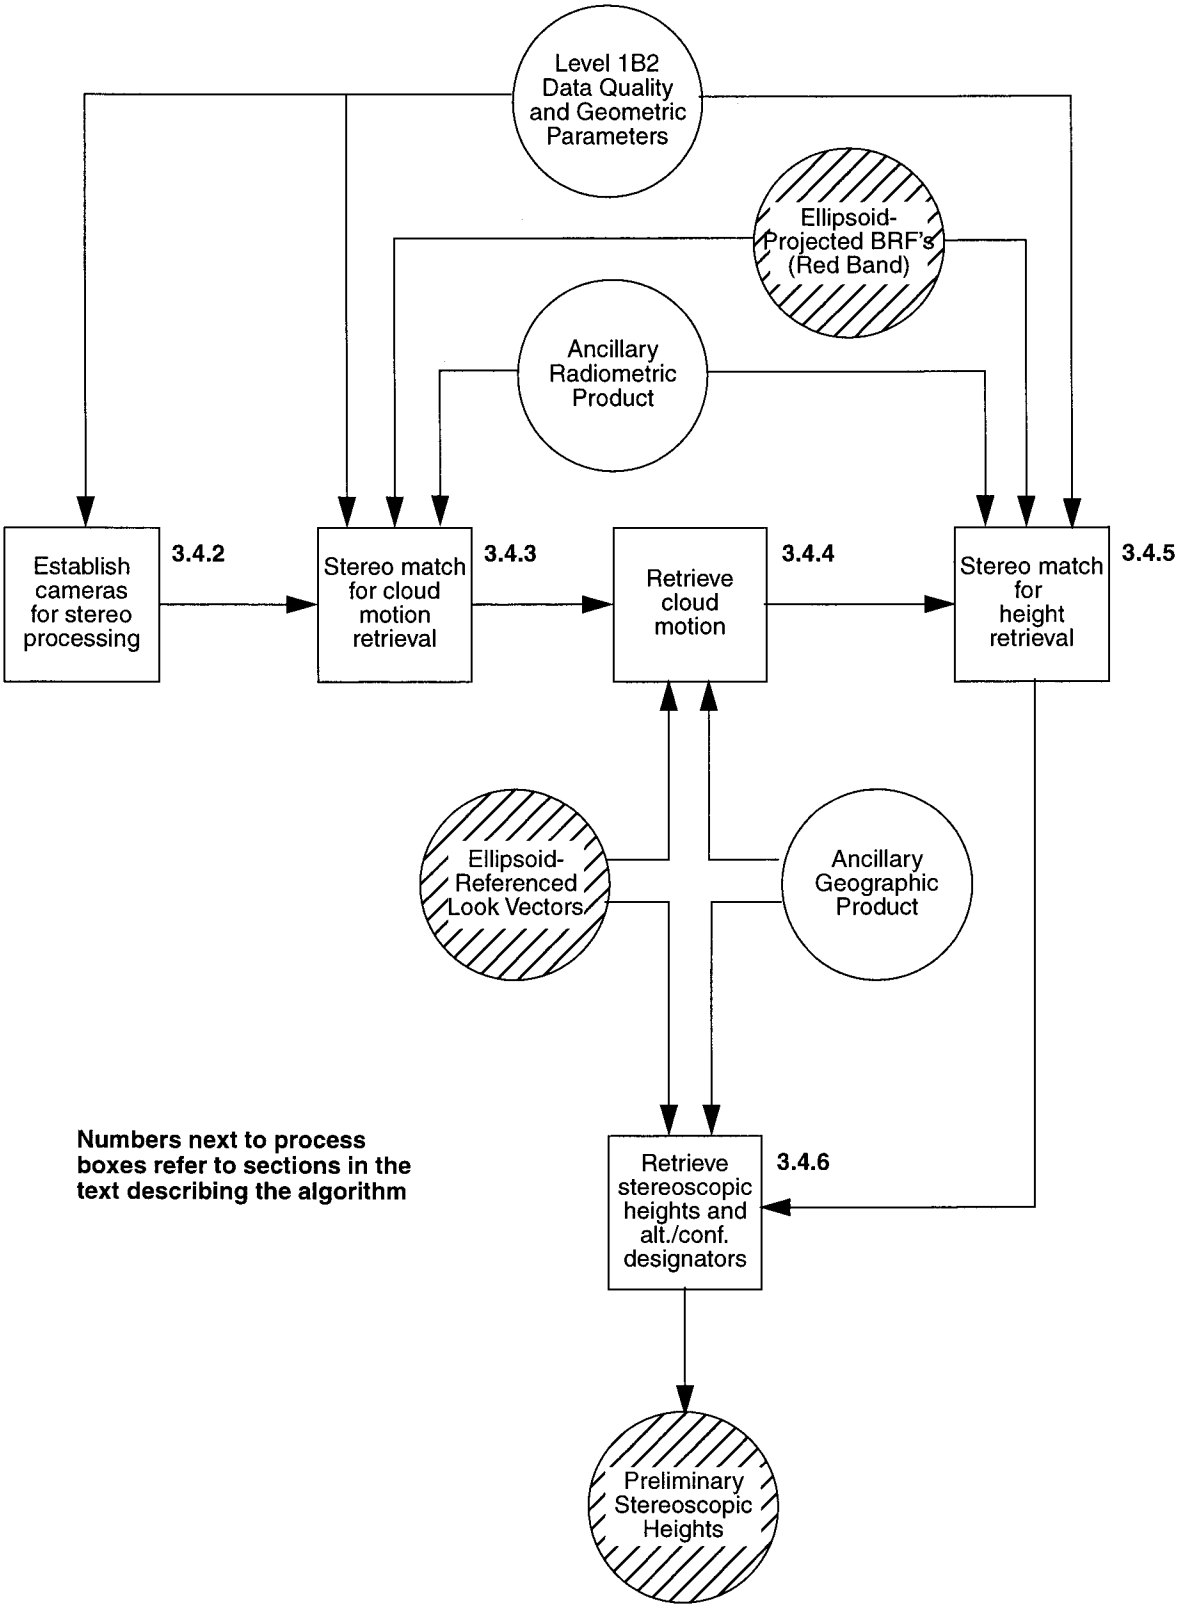

Figure 3. Stage 2 (Steps 2a - 2c) TOA/Cloud processing overview

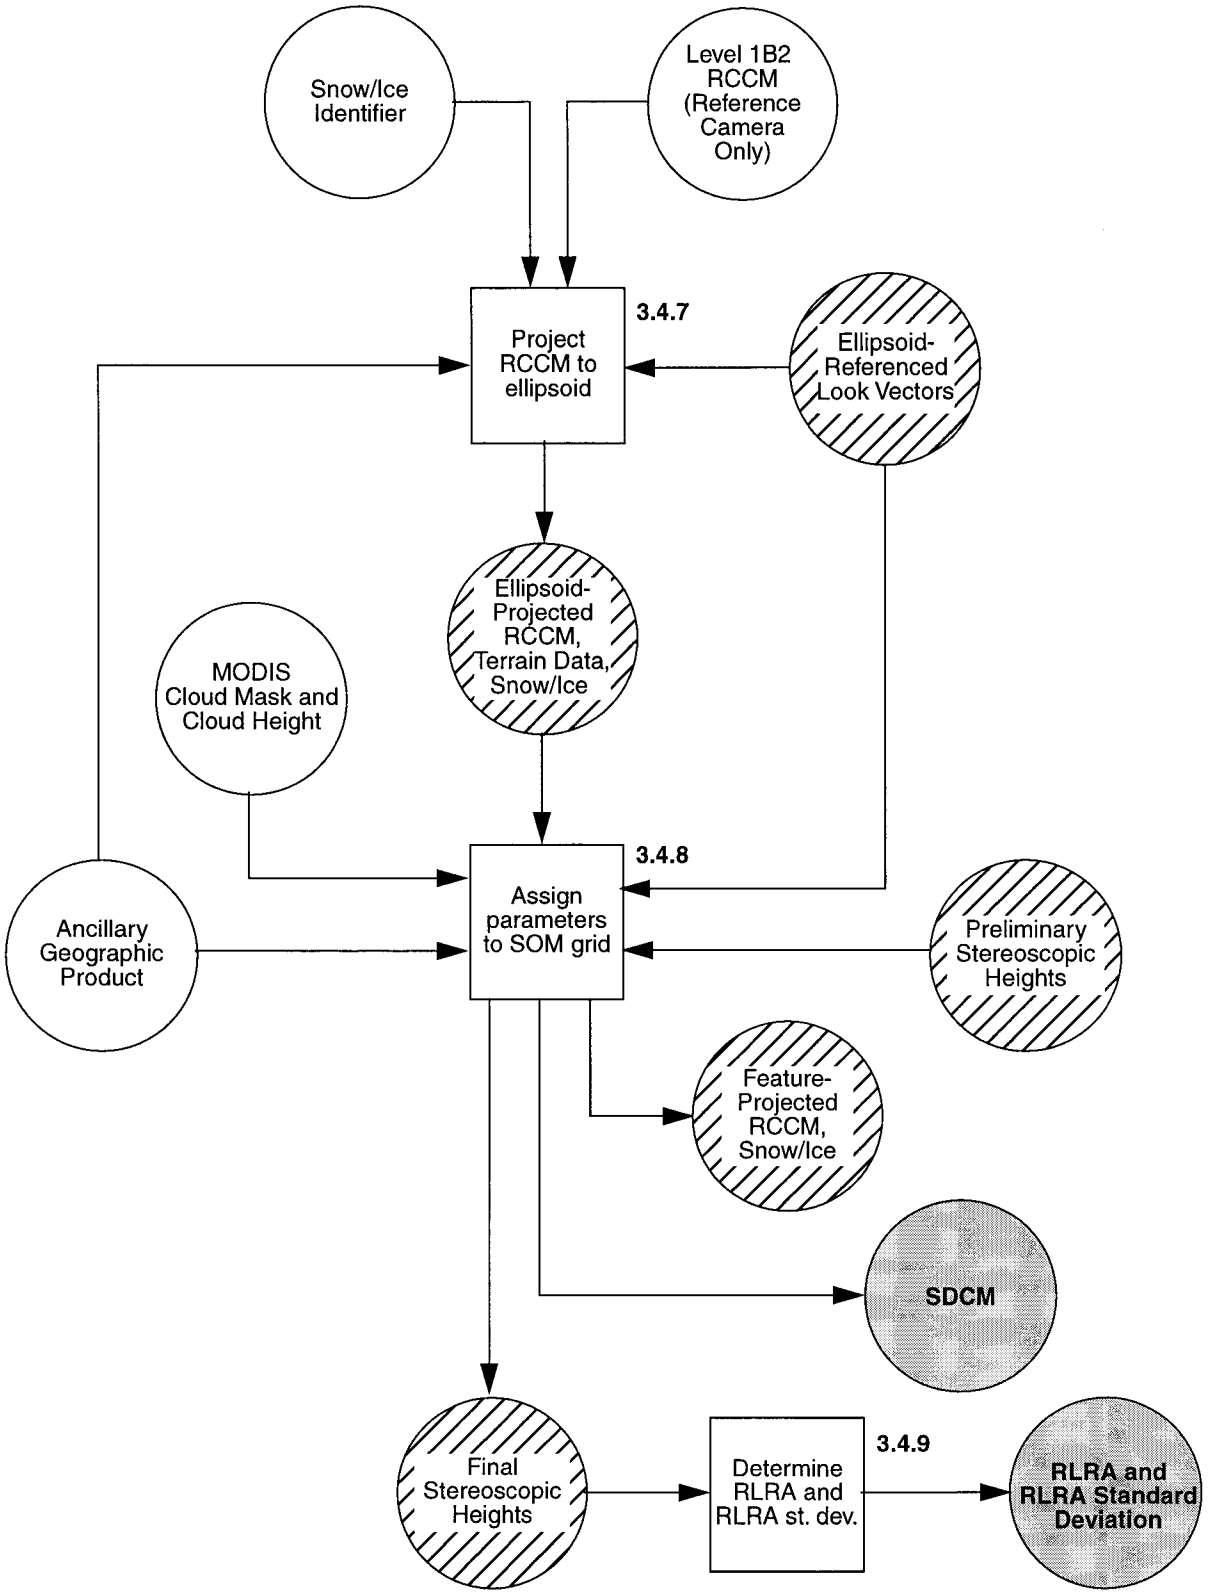

Figure 4. Stage 2 (Steps 2d - 2e) TOA/Cloud processing overview

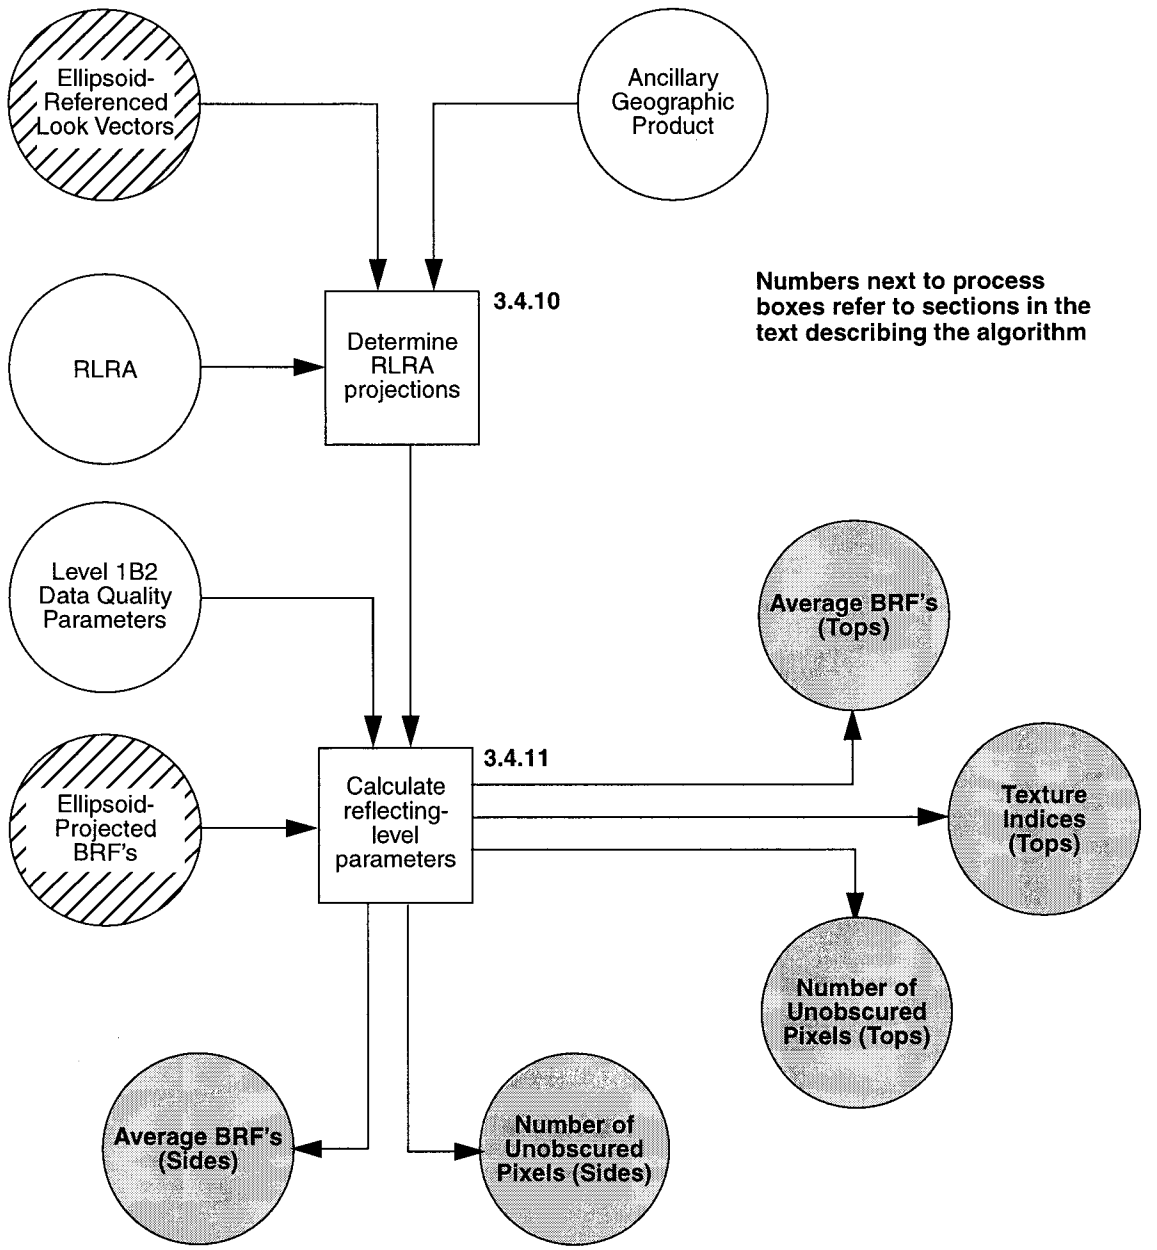

Figure 5. Stage 2 (Steps 2f - 2g) TOA/Cloud processing overview

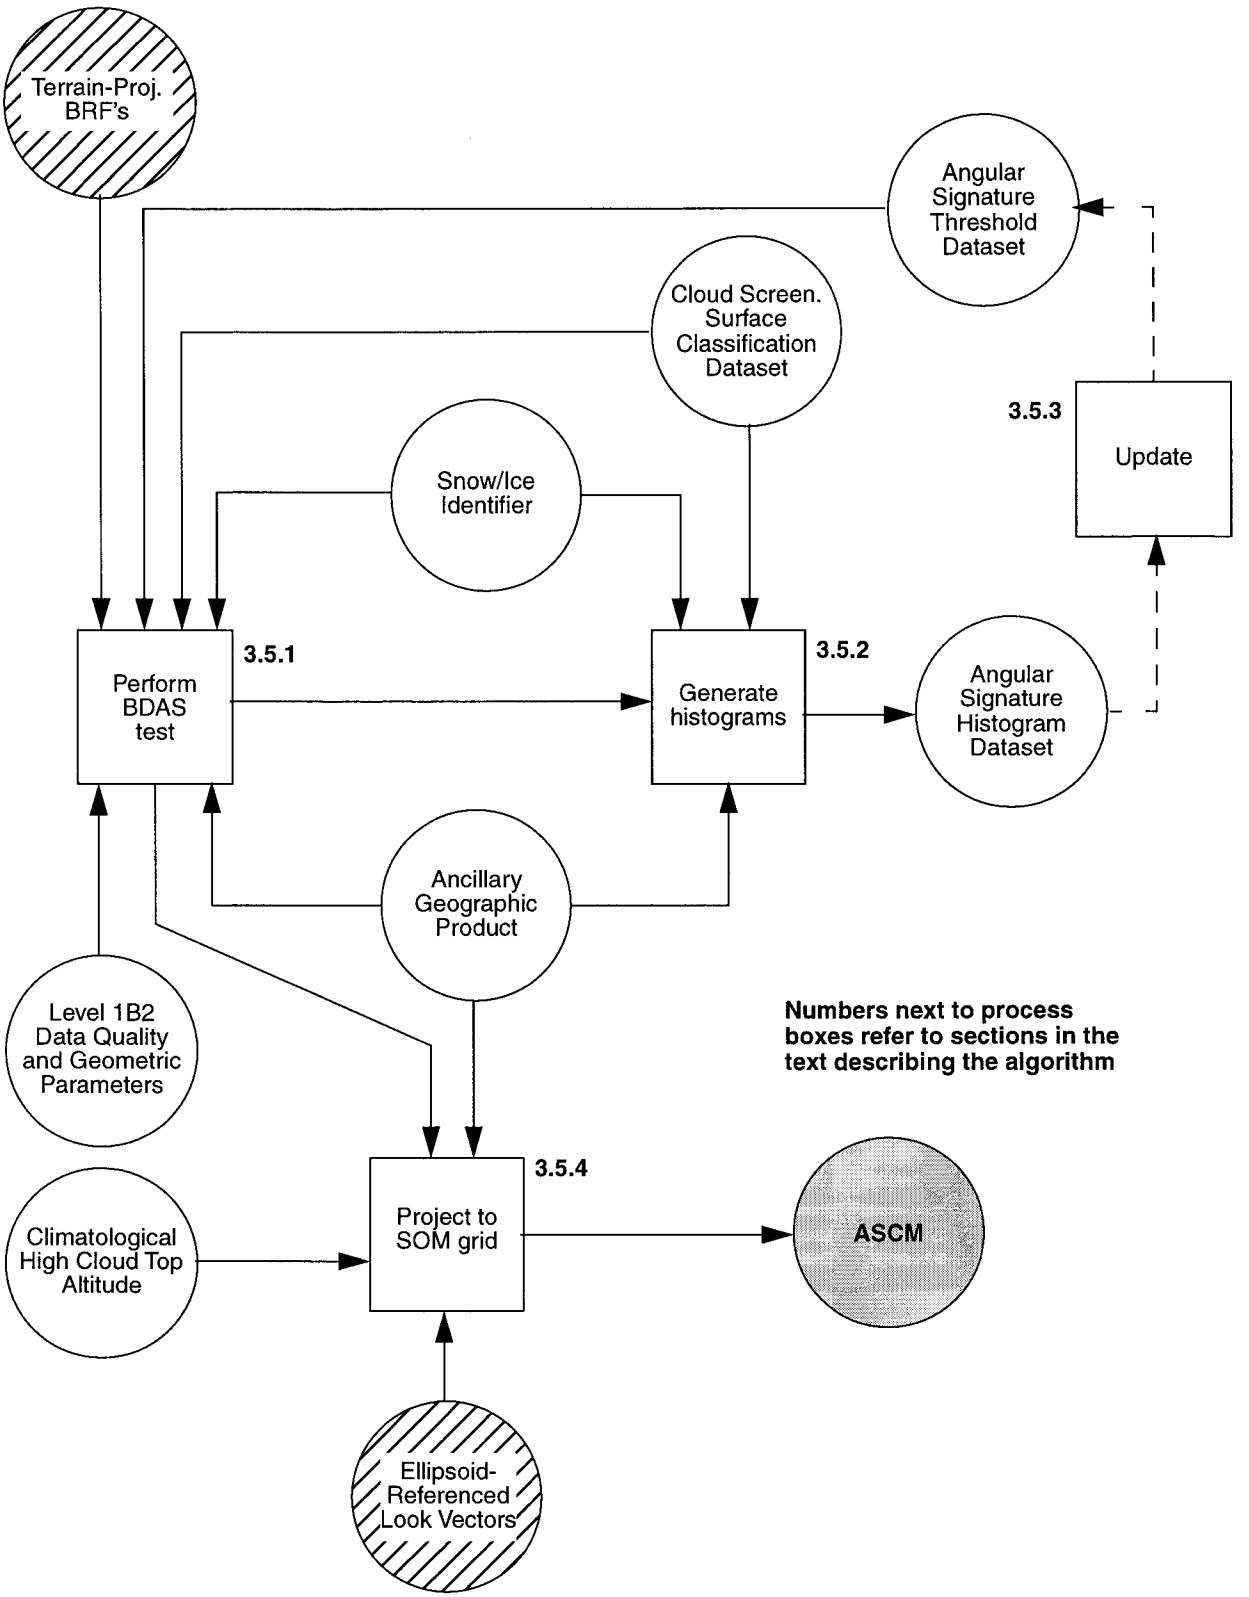

Figure 6. Stage 3 (Steps 3a - 3b) TOA/Cloud processing overview

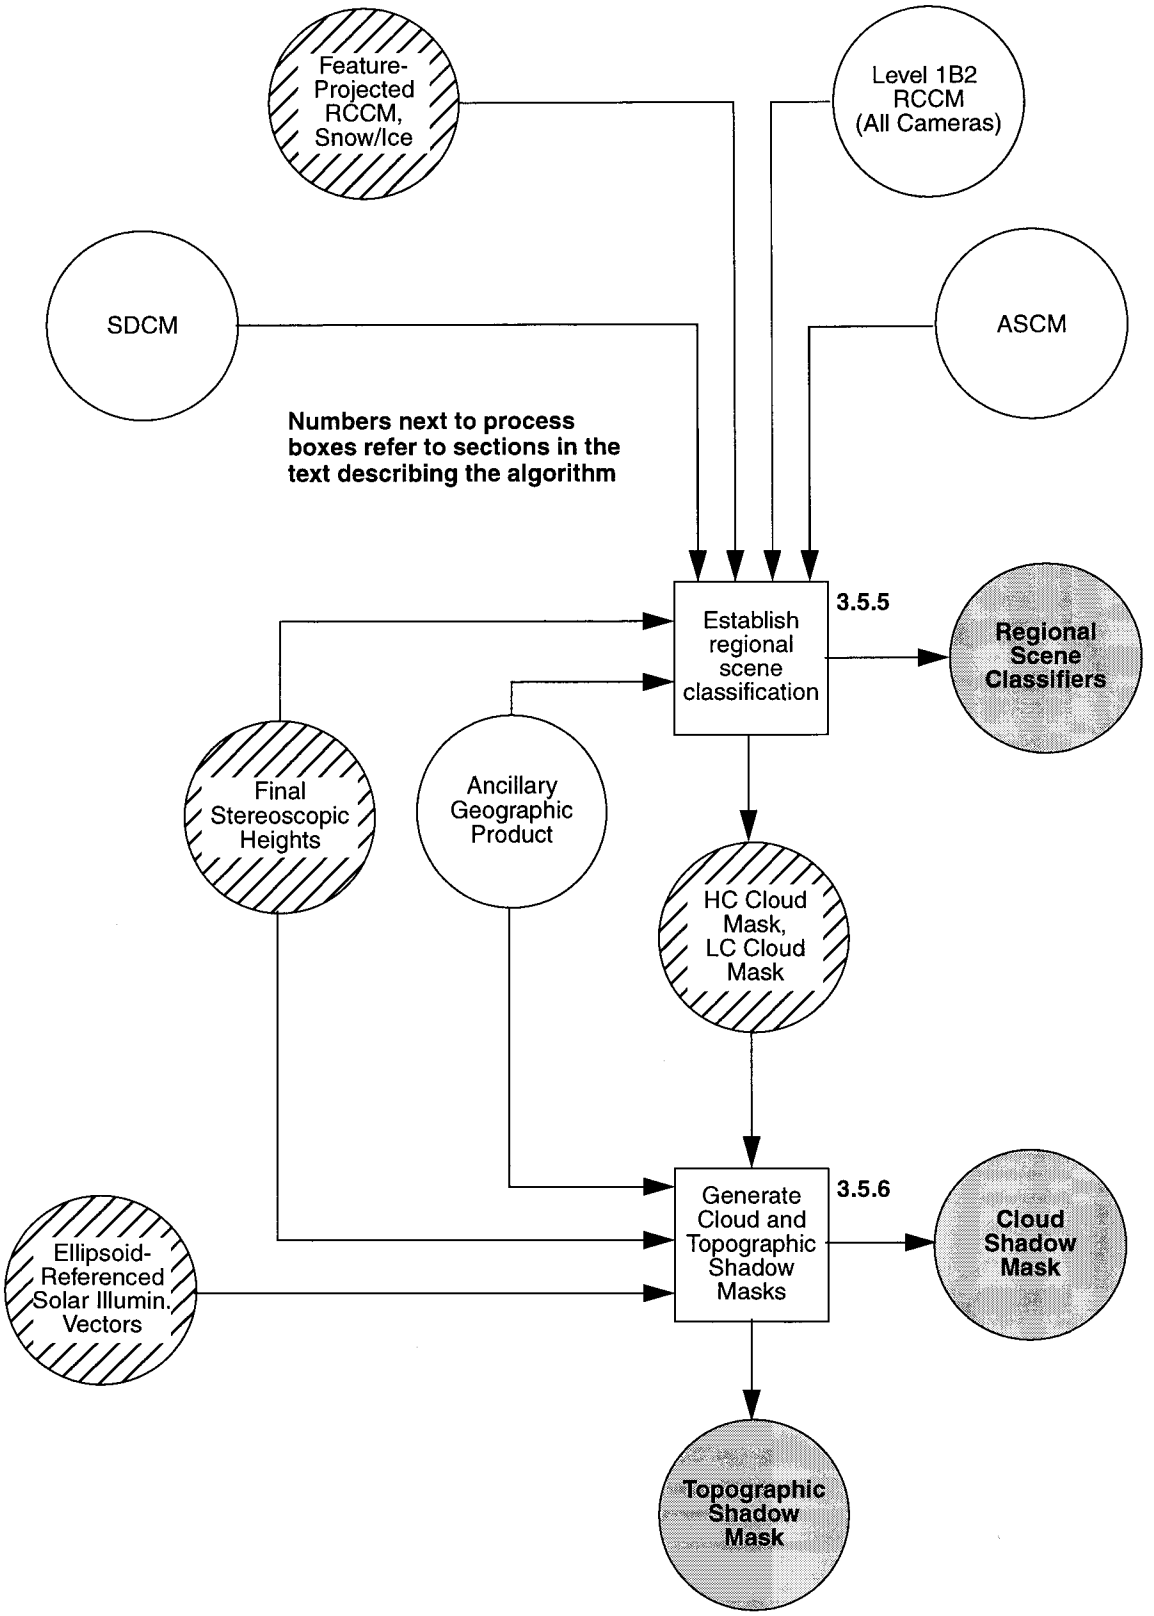

Figure 7. Stage 3 (Steps 3c - 3e) TOA/Cloud processing overview

3.2 ALGORITHM INPUT

The required inputs for cloud classification come from MISR and non-MISR sources and are summarized individually in the following paragraphs. The MISR data either come directly from the MISR processing stream, or consist of relatively static inputs, generated pre-flight by the Science Team. The latter may (e.g., see §3.2.1.13) be updated based on MISR data acquired during the course of the mission.

3.2.1 MISR data

Required inputs for cloud classification to be obtained from the MISR instrument, data system, or from the MISR team are summarized in Table 2. Further information on each of the inputs is provided below.

Table 2: Level 2 TOA/Cloud Product Inputs (MISR Data)

Input data	Source of data	Reference
Ellipsoid-referenced geometric parameters	Level 1B2 Georectified Radiance Product	*[M-4]*
Ellipsoid-projected TOA radiances	Level 1B2 Georectified Radiance Product	*[M-4]*
Terrain-projected TOA radiances	Level 1B2 Georectified Radiance Product	*[M-4]*
Data Quality Indicators and Data Flags	Level 1B2 Georectified Radiance Product	*[M-4]*
Time tags	Level 1B2 Georectified Radiance Product	*[M-4]*
Radiometric Camera-by-camera Cloud Mask	Level 1B2 Georectified Radiance Product	*[M-5]*
MISR groundtrack azimuth	Level 1B2 Georectified Radiance Product	*[M-5]*
Instrument signal-to-noise ratios	Ancillary Radiometric Product	*[M-6]*
Band-weighted exo-atmospheric solar spectral irradiances	Ancillary Radiometric Product	*[M-6]*
Land/water mask identifier	Ancillary Geographic Product	*[M-7]*
Latitude, longitude	Ancillary Geographic Product	*[M-7]*
Surface elevation	Ancillary Geographic Product	*[M-7]*
Angular signature thresholds	Angular Signature Threshold Dataset	*This document*

3.2.1.1 Ellipsoid-referenced geometric parameters

These include illumination and view zenith and azimuth angles relative to the surface normal of the World Geodetic System 1984 (WGS84) reference ellipsoid. Azimuth angles are referenced to local North. These inputs are obtained from the Level 1B2 Georectified Radiance Product.

3.2.1.2 Ellipsoid-projected TOA radiance

The ellipsoid-projected TOA radiance parameter is derived at Level 1B2 and consists of geolocated and registered radiances in all 36 channels of the instrument projected onto the WGS84 ellipsoid. A resampling process is required in order to implement this projection. The data are input at their resampled resolution, which is a function of the averaging mode at which the data were acquired, and therefore include both high resolution (unaveraged) and averaged (e.g., 4 sample x 4 line) inputs. In Level 1B2, these radiances are stored as 16-bit integer words. A multiplicative scale factor, obtained from the Ancillary Radiometric Product (ARP), is applied to convert these integers to floating point radiances, with units of W m^{-2} sr^{-1} μm^{-1}.

3.2.1.3 Terrain-projected TOA radiance

Terrain-projected TOA radiances, encoded as 16-bit integers, are input from Level 1B2 at their resampled resolution. The surface projections take into account the effects of terrain relief. Again, both averaged and unaveraged data are used. As with the ellipsoid-projected radiances, the data from Level 1B2 are multiplied by the ARP scale factor to convert the integer words to radiances in W m^{-2} sr^{-1} μm^{-1}. If the land-water mask indicates that we are over ocean, the ellipsoid-projected radiances are used in place of the terrain-projected ones. Note that these are used for the purpose of generating the Angular Signature Cloud Mask, which is based on the B, C and D cameras only.

3.2.1.4 Data Quality Indicators and Data Flags

A Radiometric Data Quality Indicator (RDQI) will be associated with each projected radiance provided by Level 1B2. This indicator will provide a representation of the radiometric quality of the input radiances used to generate values reported at in the Georectified Radiance Product. Because of the data resampling required at Level 1B2, each projected radiance represents a bilinear interpolation of four surrounding radiances obtained from the MISR images. The radiances in the imagery will be coded with a quality indicator specifying the reliability level of the radiometry on a pixel-by-pixel basis. From these, a scaled value will be produced at Level 1B2. The RDQI's take on values of 0 - 3, as follows:

RDQI = 0: Radiometric accuracy meets all specifications

RDQI = 1: Radiometric accuracy is sufficient for certain applications but some specifications are violated (see *[M-3]* and *[M-4]*)

RDQI = 2: Radiance value is available but of insufficient accuracy to be used in Level 2 retrievals

RDQI = 3: Radiance value is unavailable.

Thus, higher quality data are associated with smaller values of RDQI.

In addition to the RDQI's, radiances reported in Level 1B2 will be encoded to provide Data Flag information, for example, to indicate that a particular point on the Space Oblique Mercator (SOM) grid was topographically obscured from view by a particular camera.

Finally, Level 1B2 will provide a Geometric Data Quality Indicator (GDQI). The GDQI will provide a measure of how much image matching was used to insure high-quality image registration, relative to a pure reliance on spacecraft-supplied navigation.

3.2.1.5 Time tags

These are provided by Level 1B2 processing, and serve two purposes: (1) to derive the time interval between the same cloud feature seen by different cameras for cloud motion and height retrieval, and (2) to be used as input to the calculation of ellipsoid-referenced look vectors and solar illumination vectors.

3.2.1.6 Radiometric Camera-by-camera Cloud Mask (RCCM)

A camera-by-camera cloud mask is generated as part of Level 1B2 processing (see *[M-5]*). The observables used in generation of this mask are calculated from simple arithmetic operations on bidirectional reflectance factors (BRF's) derived from the Level 1B2 radiances. The observables are r_4, the BRF in the near-IR band at 1.1 km resolution; σ_3, the standard deviation of the 4 x 4 array of red BRF's within a 1.1 km subregion; D, a parameter derived from the red and near-IR BRF's; and D-spatial variability index (DSVI), calculated from a 3 x 3 array of 1.1 km samples and assigned to the middle one. The parameters r_4 and σ_3 are used over water, and D and DSVI are used over land. The RCCM will have confidence levels associated with it: Cloud with high confidence (CloudHC), cloud with low confidence (CloudLC), clear with low confidence (ClearLC), clear with high confidence (ClearHC), or No Retrieval (NR).

The RCCM provided by Level 1B2 is projected to the surface terrain. This means that the RCCM values are assigned to SOM grid locations at which the line-of-sight look vectors intersect the terrain. As part of the processing described in this ATB, a stereoscopic reference camera (nominally, An) is chosen and the RCCM for this camera is reprojected in two ways. First, for the purpose of combining the RCCM with stereoscopic heights, the reference camera RCCM is reprojected to the surface ellipsoid. This is done because the heights are initially associated with locations in the reference camera imagery, which is ellipsoid projected. This RCCM reprojection involves determining which SOM grid point on the surface ellipsoid, rather than the surface terrain, is intersected by the look vector. When performing this reprojection, the original terrain-projected location, and an indicator of whether that terrain location is snow or ice covered, are carried along for later use. The second reprojection involves association of a height value with the RCCM, and assigning the RCCM to the SOM grid location that lies beneath the cloud or surface feature situated at this height. This feature-projected RCCM is used in regional scene classification.

Prior to any usage of the RCCM, we check the associated glitter mask. If the RCCM is Cloud HC or Cloud LC and also flagged as glitter-contaminated, it is reclassified as No Retrieval for the purposes of the algorithms described in this document. RCCM ClearLC or ClearHC designations in the presence of glitter are probably valid and are not altered.

3.2.1.7 MISR groundtrack azimuth

In performing stereo retrievals, search windows are established in the stereo comparison cameras as the area to look for features that match the data in the reference cameras. These search windows are aligned with the spacecraft ground track. Because this ground track is not exactly co-aligned with the SOM grid (a deviation of up to 8° can occur, depending on location within the swath), the orientation angle between the ground track and the SOM along-track direction must be calculated.

3.2.1.8 Instrument signal-to-noise ratios

These are used within certain stereo matching algorithms, and are obtained from the MISR Ancillary Radiometric Product *[M-6]* for a discrete set of equivalent reflectances, which are proportional to radiances.

3.2.1.9 Band-weighted exo-atmospheric solar irradiances

These are used to convert top-of-atmosphere radiances to bidirectional reflectance factors. They are obtained from the ARP. There are several types of band-weighted exo-atmospheric solar irradiances, $\mathcal{E}_0$, contained in the ARP, distinguished by how the spectrally-varying irradiances $E_{0\lambda}$ are weighted by the spectral response of the MISR cameras. During Level 1 processing of MISR data, a correction to the observed radiances is made to account for variations in the in-band spectral response from pixel-to-pixel and camera-to-camera. As a result, the conversion of radiances to bidirectional reflectance factors required for TOA/cloud processing make use of the parameters designated $\mathcal{E}_{0,b}^{std}$ in the ARP *([M-3], [M-6])* where the superscript *std* indicates weighting by a standardized spectral response over the total spectral range for which the cameras have measurable sensitivity, and the subscript b indicates that a value of this parameter is provided for each of the 4 instrument spectral bands.

3.2.1.10 Land/water mask

This is a land/ocean/inland water/ephemeral water/coastline mask obtained from the MISR Ancillary Geographic Product (AGP). The data are provided on 1.1-km centers. The AGP is generated at the MISR SCF and stored at the DAAC. Further details of the AGP are provided in *[M-7]*.

3.2.1.11 Latitude/longitude

The AGP contains the latitude and longitude for each 1.1 km grid-center on the surface ellipsoid.

3.2.1.12 Surface elevation

In addition to the land/water mask, the MISR AGP contains surface elevation and elevation standard deviation information and the orientation of the surface normal vector relative to the WGS84 ellipsoid. The data are provided on 1.1-km centers.

3.2.1.13 BDAS thresholds

The Angular Signature Threshold Dataset is initially generated at the MISR SCF for use during Level 2 TOA/Cloud processing, and contains thresholds that are required by the BDAS algorithm. Updates to this dataset occur during the mission, through an automated analysis of histograms of the BDAS observable. In contrast with the camera-by-camera cloud mask derived at Level 1, which is generated for each camera independently of the others, BDAS requires inputs from the B, C and D cameras simultaneously. For this reason, generation of the Angular Signature Cloud Mask (ASCM), based upon BDAS, is performed during Level 2 TOA/Cloud Product generation. There exist separate threshold (and histogram) datasets for each of the three possible reference-comparison camera combinations (D-C, D-B, and C-B).

3.2.2 Non-MISR data

Inputs for cloud classification to be obtained from non-MISR sources are summarized in Table 3. The MODIS inputs are not expected to be used at launch; however, they will be integrated into the processing stream post-launch, once their influence can be evaluated.

Table 3: Level 2 TOA/Cloud Product Inputs (Non-MISR Data)

Input data	Source of data
Earth-Sun ephemeris	SDP Toolkit
Surface classification	Cloud Screening Surface Classification Dataset
Snow cover and sea ice extent	MODIS Level 2, National Snow and Ice Data Center (NSIDC), Data Assimilation Office (DAO) data, or TASC Dataset
Cloud mask and cloud-top height	MODIS Level 2
Climatological high cloud-top altitude	TASC Dataset

3.2.2.1 Earth-Sun ephemeris

This is used to obtain the Earth-Sun distance, such that observed radiances can be normalized to the standard distance of 1 AU. The source is the Science Data Production (SDP) Toolkit, which is generated by the EOSDIS Core System (ECS) contractor.

3.2.2.2 Surface classification

A Cloud Screening Surface Classification (CSSC) Dataset is required for determining the appropriate thresholds for clear/cloud discrimination in generating the ASCM. Each grid element within the CSSC Dataset is derived based on the temporal and spatial properties of the surface being relatively uniform over the grid. This allows for accurate threshold determination since the variance in the surface properties over the grid element is reduced. For further information on this dataset, see *[M-5]*.

3.2.2.3 Snow cover and sea ice extent

This is used to identify locations where the RCCM generated at Level 1 may have misclassified surface as cloud. This input is also used in generation of the ASCM to determine which combination of bands to use when calculating the observables. The current assumption is that NSIDC input, using passive microwave data, will be used at launch, to be superseded in the post-launch version of the algorithm by MODIS snow cover and sea ice retrievals. If neither NSIDC nor MODIS data is available, Data Assimilation Office (DAO) data based on NOAA retrievals are used. If none of these sources are available, TASC climatological data are used as the default.

Depending on the source of snow/ice data, the input may be in the form of a mask (Snow Covered/Not Snow Covered and Ice Covered/Not Ice Covered) or in the form of snow equivalent depths and sea ice fraction. The latter are converted to masks as described in §3.4.7.1.

3.2.2.4 Cloud mask and cloud-top height

After launch, MODIS-derived cloud-top height will be used for comparison with MISR RLRA values, and may be used as an alternative to the default values used in the at-launch version of the stereoscopic algorithm (see §3.4.8). The MODIS cloud mask would be used in conjunction with the cloud-top heights to indicate the presence or absence of cloud, as derived by MODIS.

3.2.2.5 Climatological high cloud-top altitude

This is used to establish a value for the height of cirrus clouds detected using the BDAS technique, in order to project them onto the SOM grid. This information will be obtained from climatological data contained within the MISR Terrestrial Atmosphere and Surface Climatology (TASC) Dataset, which will also contain other climatological meteorological fields *[M-12]*.

3.3 THEORETICAL DESCRIPTION: STAGE 1 PROCESSING

3.3.1 Average subregion radiances

3.3.1.1 Physics of the problem

Radiances provided by the MISR Level 1B2 Geo-rectified Radiance Product are in the same averaging mode in each channel as the data were acquired on orbit. For certain channels, these data need to be averaged to coarser resolution in order to apply the TOA/Cloud retrieval algorithms.

3.3.1.2 Mathematical description of the algorithm

Averaging over all applicable samples is used as required to generate samples with the appropriate resolution for each algorithm. Thus, the output average radiance is given by:

$$L_{av} = \frac{\sum_{i,j} w(i,j)L(i,j)}{\sum_{i,j} w(i,j)} \tag{0a}$$

where $L(i, j)$ is the radiance for the $(i, j)^{\text{th}}$ sample, and the corresponding weight, $w(i, j)$, is equal to 1 if the RDQI for the sample is $\leq \text{RDQI}_1$; otherwise $w(i, j) = 0$. We set $\text{RDQI}_1 = 1$.

An RDQI value, RDQI_{av}, is also assigned to L_{av}. Its calculation takes into account the individual RDQI's as well as what proportion of the total number of 1x1 samples that make up the subregion, N, contain valid data. In generating a 1x4 (275 m cross-track x 1.1 km along-track) subregion from 1x1 data, for example, $N = 4$, and in generating a 4x4 (1.1 km cross-track x 1.1 km along-track) subregion, $N = 16$. We define

$$\text{RDQI}_{av} = nearest\ integer\left\{\frac{1}{N}\sum_{i,j}\text{RDQI}'(i,j)\right\} \tag{0b}$$

where $\text{RDQI}'(i, j) = \text{RDQI}(i, j)$ if $\text{RDQI}(i, j) \leq \text{RDQI}_1$; otherwise $\text{RDQI}'(i, j) = \text{RDQI}_2$. We set $\text{RDQI}_2 = 3$.

In the event that $\text{RDQI}_{av} = 3$, L_{av} is set to a flag value indicating "missing data". However, if $\text{RDQI}_{av} = 3$ and any of the $L(i, j)$ is a flag value indicating "topographically obscured", L_{av} is also set to the "topographically obscured" flag value.

3.3.2 Convert radiances to BRF's

3.3.2.1 Physics of the problem

The purpose of this step is to convert TOA radiances to BRF's. The conversion is the same for all averaging modes of the Level 1B2 data, and independent of whether the input radiances are projected to the surface ellipsoid or the surface terrain. Thus, these BRF's are used as intermediate datasets prior to generation of output parameters, (e.g. BRF referenced to the RLRA). A pre-calculated solar irradiance, tailored to the MISR spectral response, is used for each of the four bands. The BRF is defined as the observed radiance at a particular illumination and view geometry divided by the radiance that would be obtained under the same conditions with a target consisting of a perfect lambertian reflector. The primary process involved is normalization to the solar input.

3.3.2.2 Mathematical description of the algorithm

The standardized band-weighted solar irradiance for each of the MISR bands for an Earth-Sun distance of 1 AU, $\mathcal{E}_{0,b}^{std}$, is obtained from the MISR Ancillary Radiometric Product. Determination of BRF r in band b is then given by:

$$r(-\mu, \mu_0, \phi - \phi_0) = \frac{\pi L(-\mu, \mu_0, \phi - \phi_0) \cdot d^2}{\mu_0 \mathcal{E}_{0,b}^{std}} \qquad (1)$$

where L is the radiance provided by the Level 1B2 Georectified Radiance Product in band b; d is the distance to the Sun, in AU, at the time of observation; μ is the cosine of the view zenith angle at the top of the atmosphere (the minus sign indicates upwelling radiation); μ_0 is the cosine of the solar zenith angle; and ϕ - ϕ_0 is the azimuthal difference between the view and solar directions. Because of the division by μ_0, BRF's are not generated when $\mu_0 < 0.01$.

3.3.3 Calculate ellipsoid-referenced look vectors and solar illumination vectors

3.3.3.1 Physics of the problem

Ellipsoid-referenced look vectors corresponding to views by the red band of each of the MISR cameras and solar illumination vectors are required for the processing described in this ATB. Time tags are used in conjunction with spacecraft navigation data, solar ephemeris information, and the MISR camera model to calculate the vectors in Earth Centered Reference (ECR) coordinates.

3.3.3.2 Mathematical description of the algorithm

Time tags, provided by Level 1B2 standard processing, are associated with the surface ellip-

soid-projected radiance data for each MISR camera. For a strip of SOM-projected data, a starting time for the strip and a set of ellipsoid transform coefficients for the red band of each camera are provided. The starting time t_0 is the imaging time of the first SOM grid element $(i_0, j_0)_{som}$ in the current strip in Coordinated Universal Time (UTC). The ellipsoid transform maps each high resolution grid center $(i, j)_{som}$ into line and sample of MISR red band imagery (l, s). The imaging time of the SOM grid point $(i, j)_{som}$ is given by:

$$t_{i,j} = t_0 + 0.0408(l - l_0) \tag{2}$$

where 0.0408 second is the image line repeat time and l_0 is the image line corresponding to the grid element $(i_0, j_0)_{som}$.

Given the imaging time of each SOM grid point, the spacecraft position $P_{spacecraft}(x,y,z)$ in ECR coordinates is determined for that time using the spacecraft ephemeris provided by Product Generation System (PGS) Toolkit functions. The look vector corresponding to the SOM grid point is then determined as the normalized vector pointing from the spacecraft location $P_{spacecraft}(x,y,z)$ to the SOM grid location in ECR coordinates, $P_{som}(x,y,z)$. The solar illumination vector is also determined by implementing PGS Toolkit calls.

For the purposes of TOA/Cloud processing, look vectors are calculated on 2.2 km centers. This is sufficient sampling such that when a look vector is needed at any point on the SOM grid, the nearest available look vector is used. A similar strategy is used for the solar vectors, though it is sufficient to provide these every 17.6 km.

3.4 THEORETICAL DESCRIPTION: STAGE 2 PROCESSING

The principal objectives of this processing step are to determine the altitude of the reflecting level and the values of parameters projected to the reflecting level. We will use the contrast measured by different cameras, and a combination of feature-based and area-based matching techniques, to determine the RLRA operationally. Determining the RLRA from multi-angle views involves detecting disparities (i.e., horizontal spatial offsets in the imagery resulting from the parallax obtained by viewing at different angles).

In the following discussion, we define a Cartesian coordinate system in which we let $+x$ represent the along-track direction of flight (roughly South), and $+y$ represent the cross-track direction (roughly East).

3.4.1 Background

The disparity of an object feature is obtained through a stereo matching process of the same feature in images observed from two or more look angles. Physically, disparity is caused by a

height difference of an object feature from a fixed point and the motion of the object along the direction where the disparity is measured. In TOA/Cloud processing the input images are registered at the surface ellipsoid. Therefore, a cloud above the ellipsoid surface causes an image disparity shown in Figure 8 as $\overline{AB_1}$, where A and B_1 represent the conjugate image features of a cloud edge.

In the absence of cloud motion, the cloud height can be obtained from intersecting conjugate look rays $\hat{a}$ and $\vec{b_1}$. However, if the cloud also has a velocity component along the direction of view, then the disparity $\overline{AB_2}$ is caused by the combined effect from both the cloud height and the cloud motion. Directly intersecting look rays $\hat{a}$ and $\vec{b_2}$ would result in an incorrect cloud height h'. In order to obtain accurate values of RLRA, it is necessary to separate the disparities due to cloud advection from those due to cloud height.

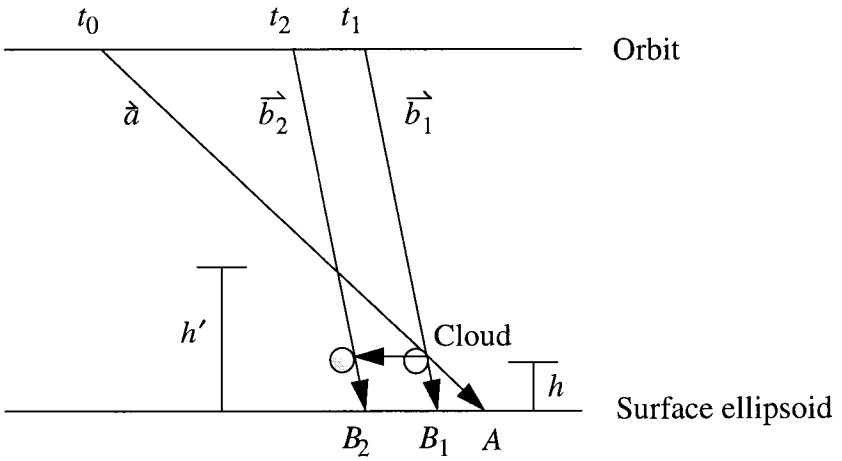

Figure 8. Disparity of cloud edge in two images with different view zenith angles

We assume that vertical cloud motion can be ignored during the 7 minute time interval between camera Df and Da views, and that the horizontal cloud motion at a given altitude is constant during this period of time over a mesoscale domain size of $(70.4\ \text{km})^2$. Then, an accurate stereo matching algorithm is used to retrieve velocity and height values for a sparse subset of features in the MISR imagery. By selecting a set of cameras with non-symmetric and oblique view zenith angles, we are able to separate the heights and velocities in the data reduction. The algorithm is first applied to a triplet of cameras (nominally, An, Bf, and Df) to detect the image disparities. The cloud height and motion of each point from this subset is simultaneously retrieved by intersecting the conjugate look rays obtained from image matching. Finally, the resulting cloud motions are binned by altitude. A similar algorithm is applied independently to a different triplet (nominally, An, Ba, and Da). The results are combined to generate the final altitude-binned cloud motion field.

Once the motion field has been derived, the conjugate points of stereo matchers applied to high resolution data from two pairs of cameras (nominally, An-Af and An-Aa) are used to retrieve

a height field, along with confidence designators from the stereo matching processing. Combining the retrieved heights with a Radiometric Camera-by-camera Cloud Mask (RCCM) from Level 1B2 results in the RLRA on which the reflecting level at 2.2 km resolution is defined.

3.4.1.1 Separation of cloud motion and height

This section studies the separability of cloud height and cloud motion. Figure 9 shows the geometry of MISR imaging in the along-track direction where cloud height and motion is highly correlated due to large view zenith angle differences from one camera to another. The shaded circles indicate the locations of a cloud at different times t_i, assuming there is a constant along-track cloud motion v_c. The projections of the cloud on the surface ellipsoid are at the locations x_i at these times, with the discrete MISR camera view zenith angles θ_i. For simplicity in this discussion, the cloud motion and the camera look vectors are assumed to lie in the along-track plane. The actual algorithms employed, described in §3.4.4 and §3.4.6, take into account the cross-track component of both the camera look vectors and the cloud motion.

If a cloud edge is seen by two cameras with different view zenith angles at times t_1 and t_2, respectively, then the traveling distance of the spacecraft and that of the cloud during this time interval can be described by the following equations:

$$v_s(t_2 - t_1) = (R + H)(\alpha_2 - \alpha_1) - (R + H)(\Gamma_2 - \Gamma_1) \tag{3}$$

$$v_c(t_2 - t_1) = (R + h)(\alpha_2 - \alpha_1) - (R + h)(\gamma_2 - \gamma_1) \tag{4}$$

where v_s and v_c are the velocities of the spacecraft and the cloud in the along-track direction, respectively, R is the radius of the Earth, H is the orbit altitude above the Earth's surface, and h is the cloud height. As shown in Figure 9, α_1 and α_2 are the angles between the initial radial line at the time t_0 and the radial line passing the image locations x_1 and x_2, respectively; Γ_1 and Γ_2 are the angles between the radial lines to the spacecraft and the corresponding image locations x_1 and x_2; and γ_1 and γ_2 are the angles between the radial lines to the cloud and the corresponding image locations x_1 and x_2.

Since $h << R$, Eq. (4) can be rewritten with variables that we are interested in:

$$v_c(t_2 - t_1) = (x_2 - x_1) + h(\tan\theta_1 - \tan\theta_2) \tag{5}$$

With multi-angle images, Eq. (5) can be generalized into a linear system as follows:

$$\begin{gathered} v_c(t_j - t_i) - h(\tan\theta_i - \tan\theta_j) = (x_j - x_i) \\ (i, j = 1, 2, \ldots, n, i \neq j, n \geq 3) \end{gathered} \tag{6}$$

The linear system expressed by Eq. (6) represents a straight line in the Δx versus Δt space in which each matching pair contributes a point to this line, where v_c and h can be determined from the slope and intercept of the line. Matching at least three images with different view zenith angles θ_i are required to solve for v_c and h. In addition, any two linear equations are dependent on each other if their determinant is zero, i.e., if

$$detA = (t_i - t_{i-1})(\tan\theta_i - \tan\theta_{i+1}) - (t_{i+1} - t_i)(\tan\theta_{i-1} - \tan\theta_i) = 0 \qquad (7)$$

In such a case, v_c and h are inseparable from these two equations. On a small scale where the flight line is straight and the surface is a plane, this singularity will always be true. On an orbital scale with an elliptic orbit and ellipsoidal surface, the traveling time interval between a pair of cameras does not linearly relate to the tangent of view zenith angles as depicted in Figure 9 and represented in Eq. (3), as long as the selected camera view zenith angles do not happen to be symmetric to the nadir view.

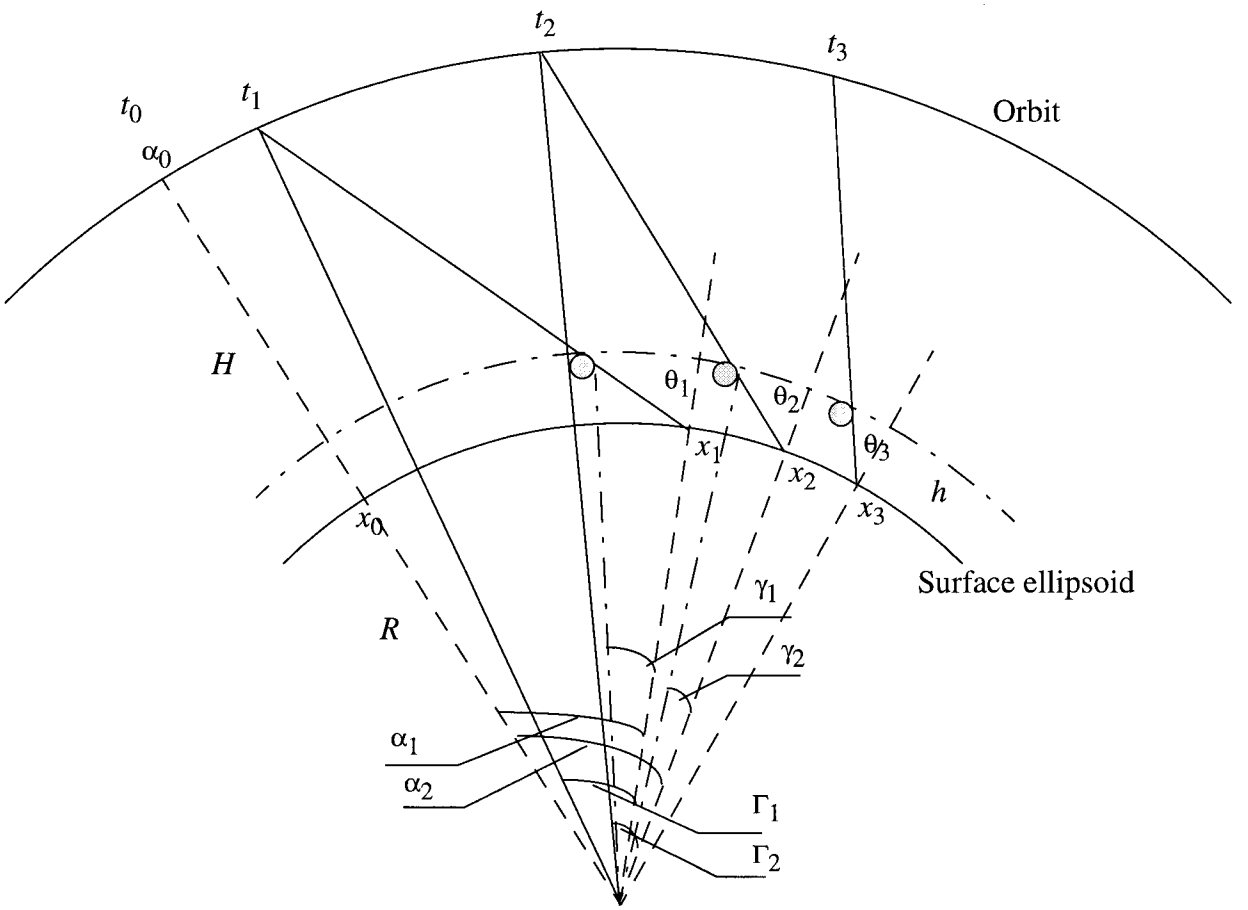

Figure 9. Geometry of MISR imaging of cloud in the along-track direction

To support this argument, Table 4 shows the relationship of MISR camera view zenith angles with respect to spacecraft traveling times (represented as line numbers) with all cameras looking

at the same point on the surface ellipsoid. The values shown in Table 4 for view zenith angles and line times are typical. They vary slightly with orbital location due to the ellipticity of the orbit and non-sphericity of the Earth. This does not change our conclusions, however. The use of typical values is sufficient for the purposes of the processing described in this document. The values shown in the table will be updated when a geometrically calibrated camera model is available. Image line number is used instead of time in Table 4 as they are linearly related. The last column contains the determinants of three indicated cameras in square brackets (e.g., [234] means Cf-Bf-Af, etc.) according to Eq. (7). The bold values indicate cases where non-consecutive cameras are used. The existence of cloud height and motion may cause the time column vary somewhat, but will not change the result. This table shows that using non-symmetric cameras, including the most oblique ones, is the best way to avoid singularity.

Table 4: Determinants of different cameras

Camera	θ_i (deg.)	$\tan\theta_i$	t_i *(line)*	$\Delta t_{i(i-1)}$ *(line)*	$\Delta\tan\theta_{i(i+1)}$	*detA (line)*
1. Df	70.5	2.82	0	_	1.09	_
2. Cf	60.0	1.73	1468	1468	0.71	-366 [123]
3. Bf	45.6	1.02	2760	1292	0.53	-115 [234]
4. Af	26.1	0.49	3887	1127	0.49	-38 [345]
5. An	0.0	0.0	5000	1113	0.49	0 [456]
6. Aa	-26.1	-0.49	6113	1113	0.53	38 [567]
7. Ba	-45.6	-1.02	7240	1127	0.71	115 [678]
8. Ca	-60.0	-1.73	8532	1292	1.09	**273 [578]**
9. Da	-70.5	-2.82	10000	1468	_	**1217 [579]**

The traveling time between a Df camera view to an An camera view is about 3.4 minutes. Assuming a cloud speed of 20 m/sec, it would contribute 4100 m or 15 lines of disparity during this interval, if the motion is entirely along-track. The longer the time span, the higher the accuracy of the cloud motion retrieval. Provided the image matching works on the most oblique angle D images, the selection of the An, B, and D cameras in the last box in Table 4 provides not only the best separability of cloud motion and height parameters but also the highest accuracy in cloud motion retrieval. The reliability of the cloud motion retrieval is optimized by performing a retrieval using a forward-viewing set of cameras and combining the result with a retrieval using an aftward-viewing set. The same argument holds true for the cross track direction, though the stereo effect is much smaller due to smaller camera side look-angle differences.

3.4.1.2 Area-based stereo matchers

Area-based stereo matchers work by designating the imagery from one camera as the reference image and the other as the comparison image. Then, a small region, or target patch, is established in the reference image and a set of comparison patches within a search window is established in the comparison image. Matching then proceeds to determine if any comparison patch in the search window successfully matches the target patch. Once this is complete, a new target patch is identified, the corresponding search window and comparison patches are located, and the process repeats. This terminology is depicted in Figure 10. In applying these matchers, the center of the comparison patch is rastered through all possible locations in the search window, beginning at a corner of the window. Thus, in certain positions of the comparison patch, some pixels in the patch extend beyond the edges of the search window. Note that the search window is aligned with the spacecraft ground track, which may deviate slightly from the SOM along-track grid direction. The patch sizes are small enough to be considered rectangles within the SOM grid coordinate system.

The tilt angle of the search window is determined by looking at two points at the left and right-hand edges of the domain in question, and projecting a feature from a given height down to the surface ellipsoid for both the reference and comparison cameras. The tilt angle is calculated for each point as the ratio of the cross-track displacement between the cameras to the corresponding along-track displacement. The final tilt angle is defined as the average of these two numbers. The displacement of each across-track string of pixels in the search window from the fiducial point is determined by multiplying the distance between each line and the along-track location of the fiducial point by the tilt angle of the window.

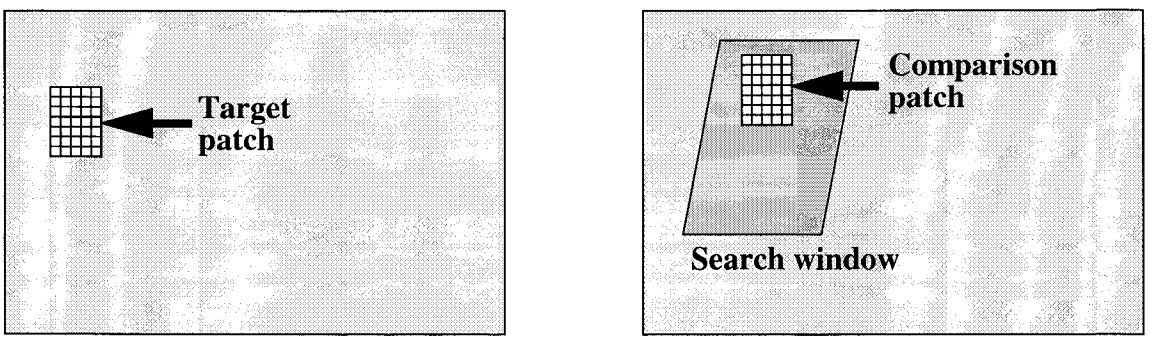

Figure 10. Terminology used for area-based stereo matchers

Application of these matchers also takes advantage of "previous match" and "pyramid" schemes to improve speed and minimize blunders. In the previous match method, the results of the previous matches immediately up and to the left of the pixel under consideration are used as the basis for a small search window. If a match that passes the appropriate threshold test is found in the course of searching this small window it is accepted and matching moves onto the next pixel.

In order for a match to serve as a valid "previous match" for this method it must pass a tighter-than-normal threshold test of its own. In the pyramid approach, the data in both the reference and comparison images are averaged to coarser resolution, and the matching is initially performed at the coarsest level. When averaging the data, a coarse pixel is only marked as having valid data if all its constituent pixels contain valid data. The best match from this level is then used as the central location for a smaller search window at the next coarsest resolution. The patch sizes always consist of the same number of averaged pixels, though they correspond to smaller and smaller physical areas as the pyramid progresses from the coarsest to the finest resolution. Matching at the last level of the pyramid is performed on the highest resolution available. MISR nominally uses a 2-level pyramid, in which the data are averaged to 550 m for the first level, and 275 m for the second level; however the number of pyramid levels is configurable.

MISR employs two area-based matchers, the Multipoint Matcher (M2) and the Multipoint Matcher using Medians (M3). The previous match method is attempted first with M2 and then with M3. If both these attempts fail, matching moves onto the image pyramid algorithm. At each level of the pyramid, M2 is applied first. If M2 fails to find a successful match, M3 is tried before moving to the next level of the pyramid. At the next finer level of the pyramid, the match from the previous level is used, regardless of which matcher found it. At the last level of the pyramid, if the final successful match resulted from M2, we also apply M3 to this single patch location to determine whether M3 verifies the point as a valid match as well. If so, the match is given higher weight in the assignment of confidence levels to the stereoscopic heights.

3.4.1.2.1 Multipoint Matcher (M2)

The Multipoint Matcher, or M2, is very accurate but limited in coverage. M2 is fast and produces the most accurate and reliable results of eleven stereo matchers tested from four different institutions with histories of research work in this area. Blunders are rare and there is a trade-off between blunders and coverage of the matcher. M2 has been applied to a synthetic stereo-pair based on AVIRIS data and to a wide variety of different cloud types (including complex multi-level clouds) in a long strip of Along-Track Scanning Radiometer (ATSR) forward-nadir stereo data at 1.6 μm. M2 has also been tested with simulated cloudy-sky images as well as AirMISR data. It works with arbitrarily distorted images using local affine transformations. In the case of MISR, however, the Level 1B2 images are used as input and these have most of the distortions removed prior to matching. The algorithm is applied to determine the conjugate locations in the high resolution (i.e., unaveraged) BRF imagery, thus pinpointing the disparities accurately. For a stereo image pair, M2 takes the target patch in the reference image and a set of comparison patches within a search window in the comparison image and computes a matching metric. This metric is computed by taking all the BRF values in each patch, subtracting the mean BRF within the patch from each pixel, then normalizing by the difference in the maximum and minimum BRF's. Then, the absolute difference between these values in the target patch and the corresponding values in the comparison patch, summed over the area of the patches and normalized by an uncertainty estimate,

is tested against a threshold. The M2 metric is defined as follows:

$$S_{M2} = \frac{\sum_{i,j} \left| \left[\frac{R(x_i, y_j) - \langle R \rangle}{R_{max} - R_{min}} \right] - \left[\frac{C(x_i, y_j) - \langle C \rangle}{C_{max} - C_{min}} \right] \right|}{\sigma_{M2}} \tag{8}$$

where $R(x_i,y_j)$ is the BRF in the reference image pixel value at (i, j), and $C(x_i,y_j)$ is the corresponding BRF in the comparison image. In order to apply this matcher we require that every value of $R(x_i,y_j)$ and $C(x_i,y_j)$ within their respective patches be associated with a Radiometric Data Quality Indicator (RDQI) $\leq$ $RDQI_3$. We set $RDQI_3 = 1$. Thus, a single "bad" pixel in either patch will cause matching to be skipped. In the notation used in Eq. (8), i and j are relative indices within the patches for summation, but the absolute values of the pixel coordinates within the reference and comparison images are not necessarily the same. In Eq. (8), $\langle R \rangle$ and $\langle C \rangle$ are the average BRF values within the target and comparison patches, respectively; R_{max} and R_{min} are the maximum and minimum BRF values within the target patch, respectively; and C_{max} and C_{min} are the maximum and minimum BRF values within the comparison patch, respectively. The quantity, σ_{M2}, by which Eq. (8) is normalized, is an estimate of the uncertainty in the numerator of Eq. (8), and is defined below. For later use, we also define

$$N_{pts} = \sum_{i,j} 1 \tag{9}$$

i.e., it is the number of pixels in the patch.

We first note that the uncertainties in the terms in the numerator of Eq. (8) are independent of the absolute radiometric accuracy of the instrument, the band-to-band uncertainty (because only the red band data are used), or the camera-to-camera uncertainty (because any unknown scaling error cancels out as a result of the form of the expression). Additionally, we do not include systematic pixel-to-pixel uncertainties as the patches are localized with the cameras' fields-of-view. Thus the principal contributors to σ_{M2} are random pixel-to-pixel uncertainties and the effects of any subpixel misregistration between the patches, the latter expected to be more significant due to the high signal-to-noise ratio of the MISR cameras. We estimate the value of σ_{M2} by setting it equal to the value that the numerator of Eq. (8) would take if a target patch were being matched to itself, but with a misregistration in the x or y directions. We assume that the mean, maximum, and minimum BRF of the misregistered patch is the same as the target patch, and that the value of the misregistered patch BRF $R_{misreg}(x_i,y_j)$ is estimated as

$$R_{misreg}(x_i, y_j) = \frac{R(x_{i-1}, y_j) + R(x_{i+1}, y_j) + R(x_i, y_{j-1}) + R(x_i, y_{j+1})}{4} \tag{10}$$

i.e., the average of the BRF's at the four nearest locations, since we do not know in which direction the hypothetical misregistration occurs. We then make the approximation that this average may be estimated by the average BRF in the whole patch, $\langle R \rangle$. Using these arguments, we derive

$$\sigma_{M2} = \sum_{i,j} \left| \left[\frac{R(x_i, y_j) - \langle R \rangle}{R_{max} - R_{min}} \right] \right| \tag{11}$$

Note that for a patch which is nearly uniform in BRF, the effect of misregistration is not so significant and σ_{M2} merely provides an estimate of the random noise on the data. As the contrast in the patch increases, σ_{M2} increases in magnitude, since the effect of misregistration becomes more significant for patches with spatially varying BRF. The inclusion of σ_{M2} as a normalization factor in Eq. (8) means that a value of S_{M2} on the order of unity corresponds to a match to within the capability of the instrument to detect a difference in the patches.

For a given target, the candidate x and y values of disparity are those for which S_{M2} is smaller than or equal to a threshold T_{M2}, which is established to remove any possible blunder resulting from the matching. If $S_{M2} > T_{M2}$, the disparity is deemed an unsuccessful match and is discarded.

Ideally, the result of the threshold test will eliminate all candidate disparities except one. If a unique disparity (x and y value) is obtained, the match is considered successful. The match is considered to have failed if either (a) all candidate disparities are eliminated by the threshold test, or (b) there remain multiple candidate disparities following application of the threshold test, and an unambiguous "best" answer cannot be chosen.

The ambiguity test works as follows. If multiple matches from M2 satisfy the threshold criterion, take the "best" match to be the one that minimizes S_{M2}. Let $S_{M2,min}$ be the corresponding value of the matcher metric. Now, set a secondary threshold A_{M2}, given by $A_{M2} = f \cdot S_{M2,min}$, where f is a value slightly larger than unity. We set $f = 1.1$. If the best match is the only one for which $S_{M2} < A_{M2}$, it is considered an unambiguous successful match. If there are still multiple matches below the threshold A_{M2}, we apply a spatial cluster test to these matches. To implement the cluster test, we compute two parameters: s_x, the maximum absolute value of the difference between all x disparities that passed both the primary ($S_{M2} \leq T_{M2}$) and secondary threshold tests, and s_y, an analogously defined parameter for the y disparities. If $s_x \leq$ a threshold value p_x, and simultaneously $s_y \leq$ a threshold value p_y, the disparities are said to be spatially clustered. At present, we set $p_x = 3$ pixels and $p_y = 3$ pixels. If the disparities satisfy the clustering criteria, the best match is taken as the successful result. Otherwise, we deem the results as ambiguous and consider there to be no successful match. The ambiguity test is used after application of the "previous-match" method and at all but the last level (highest resolution) of the pyramid scheme.

We note that if $R_{max} - R_{min} = 0$ or $C_{max} - C_{min} = 0$, the M2 matcher cannot be applied.

3.4.1.2.2 *Multipoint Matcher using Medians (M3)*

A second area-matching algorithm, called the Multipoint Matcher using Medians, (M3) is also used. M3 works similarly to M2, and has been tested against simulated cloudy-sky data, AirMISR data, and AVHRR data that have been artificially enhanced to MISR resolution. It uses medians rather than means, and applies a normalization scheme that yields better coverage than M2, with minor loss of accuracy. M3 is applied in a similar manner as M2, except that the matching metric is given by:

$$S_{M3} = \frac{median_{ij}\left\{\left|\frac{R(x_i, y_j)}{median(R)} - \frac{C(x_i, y_j)}{median(C)}\right|\right\}}{\sigma_{M3}} \tag{12}$$

where the inclusion of σ_{M3} follows the same rationale as with M2 but is defined differently, i.e., the BRF values in each patch are normalized by dividing by the median value of each patch. A distribution of the absolute values of the differences in the corresponding pixels between the patches is then formed (i.e., the absolute value of the difference of the top right hand pixels in each patch, etc.), and the final metric is the median of this distribution, normalized by σ_{M3} as shown in Eq. (12).

The derivation of σ_{M3} uses similar methodology as in the derivation of σ_{M2}, with the exception that instead of estimating the value in Eq. (10) using the mean BRF in the patch, we use the median. With this assumption, we obtain

$$\sigma_{M3} = median_{ij}\left\{\left|\frac{R(x_i, y_j)}{median(R)} - 1\right|\right\} \tag{12a}$$

Again, for a given target patch, the candidate x and y values of disparity associated with the M3 test are those for which S_{M3} is smaller than or equal to a threshold T_{M3}. A different threshold value may be used for M3 than is used in M2. However, as with M2, we consider a successful matching to have occurred if the application of the threshold results in only a single (i.e., unique) disparity, or if the ambiguity test allows a single selection to be chosen from multiple disparities, in which case the x and y values of disparity associated with the minimum value of S_{M3} are chosen.

We note that if $median(R) = 0$ or $median(C) = 0$, the M3 matcher cannot be applied.

3.4.1.3 Feature-based stereo matchers

3.4.1.3.1 Nested Maxima (NM)

A feature-based matcher, called Nested Maxima (NM), is employed to identify and match a relatively few features very quickly, but to do so unambiguously with a high degree of confidence. The algorithm uses a simple iterative inequality logic which successively finds local maxima in the signals, and then the local maxima of the local maxima, and so on. The features are then compared for possible matches based on the range of allowable disparities. The NM algorithm has been tested on a number of stereo pairs of images (one member of the pair being high resolution radiometric data of real cloud scenes, the other being a simulated image based on a "known" disparity field). The simulated disparity field was calculated to mimic that expected for MISR observations. NM has also been tested on AirMISR data. Matches were typically found for 1% of all samples within the scene, at a 98% accuracy rate. We expect this performance to carry over to MISR data.

Local maxima are defined within a one-dimensional string of numbers with respect to a search radius, m. If m is set to 1, then a point at location i in the string is defined to be a local maximum if its BRF is larger than the BRF's at location i - 1 and i + 1. If m is set to 2, then the same point at location i will only be a local maximum if the previous condition is satisfied, and additionally, if the BRF at i - 1 is greater than the BRF at i - 2 and if the BRF at i + 1 is greater than the BRF at i + 2. The same logic continues for higher values of m.

Nested Maxima is applied separately to pairs of cameras (e.g., Bf-An and Bf-Df). The reference image of the pair is chosen to be the appropriate forward or aftward wind retrieval reference camera, (nominally Bf and Ba). The conjugate cloud features are first identified independently within one-dimensional along-track strings of data within these two images and then matched. Within these images, target and search windows are defined in order to establish the areas within which valid matches may exist. Figure 11 shows schematically how NM is applied. Note that the search window is aligned with the spacecraft ground track, whereas the strings are aligned with the SOM grid. The tilt angle of the search window relative to the SOM grid is nearly identical for the reference and comparison images.

The ellipsoid-projected BRF's within the one-dimensional strings in the reference image are first arranged in a sequence based on their grid position. Taking r_i^0, $i = 1, 2, \ldots, N^0$, as the sequence of valid BRF's (defined to be BRF's for which the RDQI $\leq$ RDQI$_3$) contained in a one-dimensional along-track string of data, the first pass identifies those indices i' which are local maxima. This forms a sequence of indices $i' = i'_1, i'_2, \ldots, N^1$ ($i'_1 \geq 1$; $N^1 < N^0$) which identifies the next sequence of BRF's r_i^1, where

$$r_i^1 = r_{i'}^0 \tag{13}$$

In practice, the algorithm simply keeps track of the relevant indices at this level. This step is repeated, and the indices stored, level by level, until the highest level is reached (i.e., when it is no longer possible to find local maxima) or until a maximum level number is reached, whichever comes first. We set the highest level to 5. At all levels of sifting except the first, the search radius, m, is set to 1. At the first level, we set $m = 2$. Sifting is performed for each adjacent one-dimensional along-track string of data within the reference image. Stacking the results from adjacent strings together generates a set of indices at each sort level for which both the x and y values are recorded. The same procedure is then repeated for the comparison image.

An additional note of practical significance is that at any level except the lowest, we allow for local maxima to be selected at the endpoints of the one-dimensional strings. For the endpoints the comparison test described above is only made relative to points that are interior to the string.

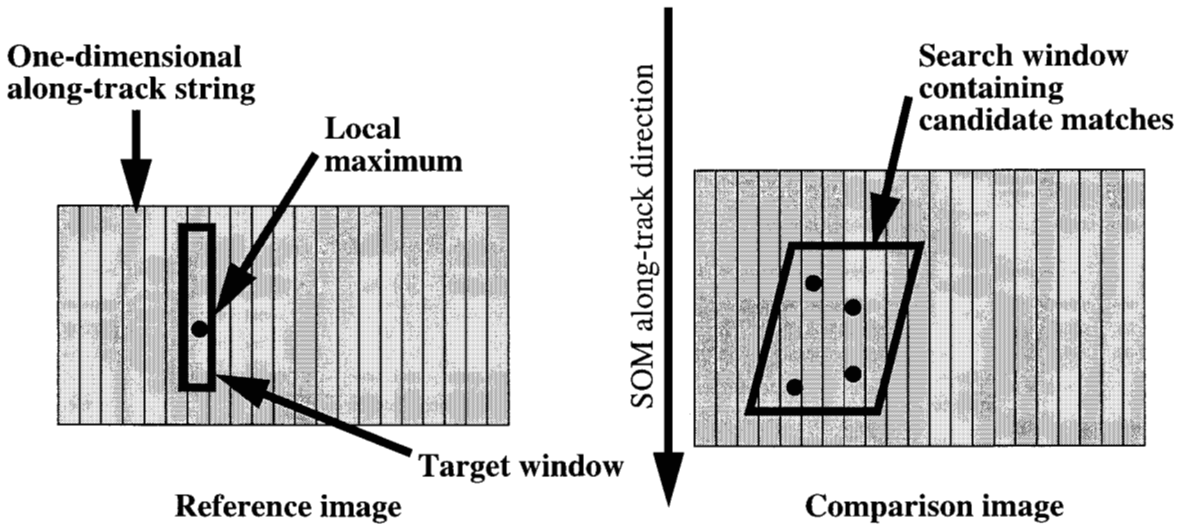

Figure 11. Terminology used in application of Nested Maxima

Once the arrays of local maxima are obtained at each level within the reference and comparison images, the following procedure is applied. At each level, beginning with the highest level number that is common to both the reference and comparison images, the located maxima are examined to isolate a smaller set of candidate matches. Stepping through each along-track string in the reference image, each local maximum is considered in turn to find those local maxima in the comparison image which are possible matches to it. For each local maximum at a given level in the reference image, the following steps are followed:

(1) Establish a search window in the comparison image, surrounding the location corresponding to the reference image local maximum being considered. The dimensions of the search window, in pixels, are $[2U_c + 1]$ x $[U_a(\text{leading}) + U_a(\text{trailing}) + 1]$, where the subscripts "$c$" and "$a$" refer to cross-track and along-track, respectively. The parameters U_a(leading), U_a(trailing), and U_c are defined below. The search window is then tilted to align with the spacecraft ground track.

(2) Step through the portions of each along-track string that are contained within the search window, and find all the unique matches. Note that the along-track strings may be shorter at the search windows' edges due to the tilt (see Figure 11). Determination of whether the uniqueness condition is satisfied involves the following steps:

(a) Establish whether only one "unused" local maximum at or above the current level is located within the one-dimensional along-track string in the search window.

(b) If (a) is satisfied, establish a target window in the reference image surrounding the location corresponding to the comparison image local maximum being considered. The along-track dimension and positioning of the target window in the reference image are established in identical fashion to the manner in which the along-track dimension and positioning of the search window in the comparison image are established, with the exception that the "leading" and "trailing" dimensions are reversed (see below). The target window is only 1 pixel wide. No tilt is applied to it.

(c) Establish whether the target window contains the currently considered reference maximum as the only unused one present within its one-dimensional along-track string.

(3) All local maxima in the comparison image which survive step (2) are declared to be the final candidate set of matches associated with the currently considered local maximum in the target window.

(4) The final match from this candidate set of matches is chosen by constructing patches around the pixels containing the local maxima and calculating the M2 metric for all possible sets of matches, and picking the one that (a) has the minimum value of this metric, (b) meets the threshold criterion, and (c) satisfies the ambiguity tests of §3.4.1.2.1. In this manner we employ the main strength of NM which uses the nesting of the maxima to quickly select a small subset of possible matches, and then the suitability of these points is determined using M2.

(5) If steps (1) - (4) result in a single local maximum in the comparison window that is matched with the currently considered point in the target window, both features are marked as "used" and removed from further consideration at this and all lower levels. In order to make the algorithm produce identical results regardless of the direction of the cross-track search, the matches are not marked as "used" until the processing of the level is completed.

Steps (1) - (5) are then repeated for all other local maxima in the target window. When this is complete, this procedure is repeated for the next level of matches. This process terminates when the second lowest level of nested maxima is reached.

The values of U_c are determined largely by displacements due to the cross-track component of the wind vector. Thus, U_c is symmetric in the cross-track direction. On the other hand, U_a is determined from a combination of the along-track component of the wind vector and the apparent displacement due to cloud height. Thus, the values of U_a must be distinguished by whether they

are "leading", i.e., further along the flight direction than the reference point, or "trailing". This requires us to position the search and target windows asymmetrically. This is illustrated in Figure 12.

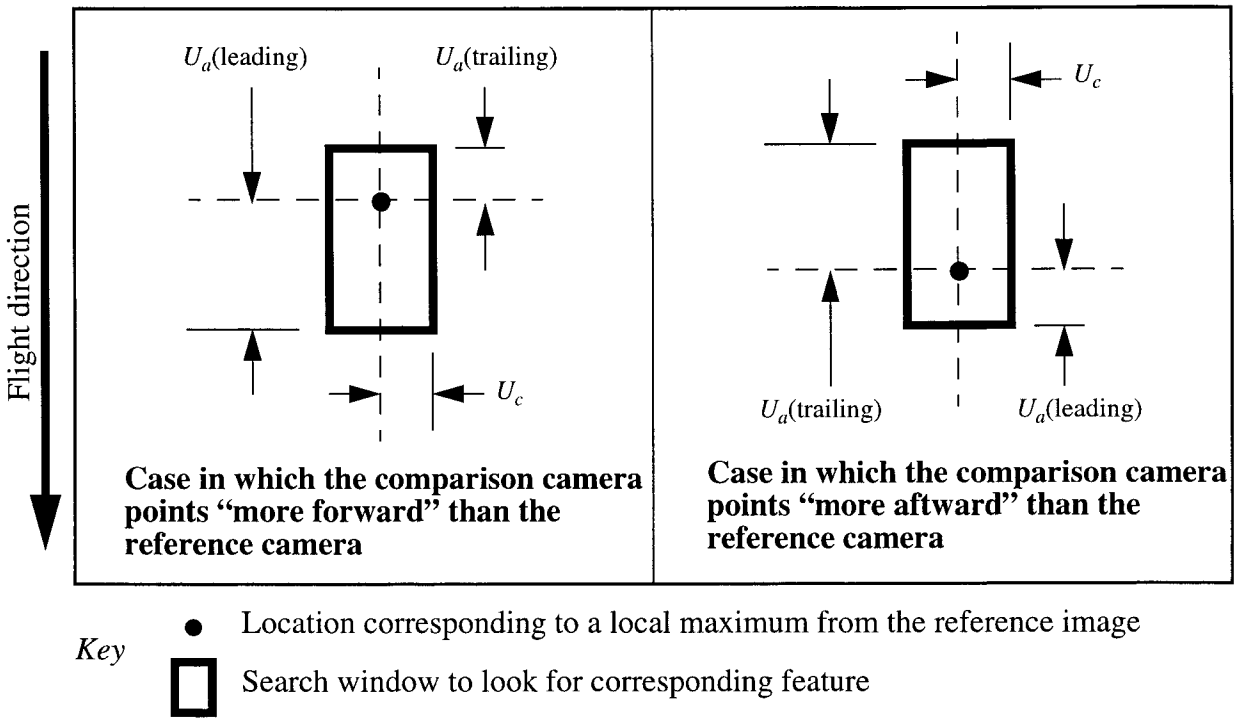

Figure 12. Dimensions of search windows used in Nested Maxima

The values of U_a(leading), U_a(trailing), and U_c are expressed as numbers of 275-m pixels, and depend on which cameras are selected as the reference, comparison, and common comparison cameras in the wind retrieval triplets. The following equations are used to determine the numerical values for a particular pair of cameras (i.e., reference-comparison or reference-common comparison), and assume a near-epipolar geometry with 1 pixel (275 m) non-epipolarity for each camera, 1 pixel (275 m) cross-track navigation uncertainty for each camera, a maximum cloud height, h_{max}, of 20,000 m, and a maximum cloud speed, v_{max}, of 100 m/sec. The extra term in the equation (l) is added to account for uncertainty in the image registration. It is currently set to 1.1 km or 4 x 275 m pixels.

$$U_c = \frac{v_{max}\Delta t}{275\ \text{m}} + l \tag{14}$$

For the case in which the comparison camera points "more forward" than the reference camera:

$$U_a(\text{trailing}) = U_c$$

$$U_a(\text{leading}) = U_c + \frac{h_{max}}{275\ \text{m}} \cdot \left|\tan\theta_2 - \tan\theta_1\right| \tag{15}$$

and for the case in which the comparison camera points "more aftward" than the reference camera:

$$U_a(\text{leading}) = U_c$$

$$U_a(\text{trailing}) = U_c + \frac{h_{max}}{275\ \text{m}} \cdot \left|\tan\theta_2 - \tan\theta_1\right| \tag{16}$$

In Eqs. (14) - (16), Δt is the absolute value of the time difference between acquisition of the imagery by the two cameras in the pair (obtained by taking the Δt values shown in Table 4, expressed as line numbers, and multiplying by 40.8 msec, the MISR camera line repeat time), θ_1 is the along-track view zenith angle of the reference camera, and θ_2 is the along-track view zenith angle of the comparison (or common comparison) camera. Note that for the above equations to give the proper results, the θ angles are specified to be positive for camera look vectors viewing forward of nadir, and negative for camera look vectors viewing aftward of nadir. Because only nominal values of the along-track view zenith angles are required, we use the values from Table 4 (i.e., predetermined values). The addition of 4 pixels in Eq. (14) is done to account for the non-epipolarity and navigation error contributions from each camera. The results of Eqs. (14) - (16) are rounded up to the nearest integer after all the calculations have been done. The round-up of U_c is performed after U_a(leading) and U_a(trailing) have already been calculated.

The algorithm for calculating the tilt of the flight direction relative to the SOM grid is as follows:

(1) Find the look vectors for the reference and comparison cameras at the middle of the far left and far right edges of each domain within which matching is performed.

(2) Determine *dx* and *dy*, the along-track and cross-track components of these vectors in the plane of the SOM grid.

(3) Determine the slope of the look vector differences, $\frac{\Delta dx}{\Delta dy}$ for the two cameras, at each domain edge.

(4) The slope of the tilt is then given the by the average of the values at the two domain edges.

In applying NM to a reference or comparison image along-track string containing N points, among which the number of valid values (i.e., those for which RDQI $\leq$ RDQI$_3$) is N^0, we require

$$\frac{N^0}{N} \geq f_{NM} \tag{17}$$

where f_{NM} is a configurable parameter that establishes how many missing data points are allowed. At launch, we plan to set f_{NM} to an initial value of 0, i.e., missing data values will not affect the calculation of local maxima. As experience is gained with actual data, this condition may be tightened. We nonetheless always require the condition that at the first level at which local maxima are established, a BRF cannot be declared a local maximum if there are any invalid data (i.e., RDQI > $RDQI_3$) within the search radius m.

3.4.1.3.2 Rank Sort (RS)

A second feature-based matching technique, known as Rank Sort (RS), is also used. RS is a one-dimensional algorithm, and it is therefore necessary to apply it to multiple along-track strings of data in order to allow for geometric uncertainties in the data due either to image navigation errors incurred during Level 1B2 processing or errors in the cloud motion retrieval. For example, the time interval between views of a feature located close to the surface ellipsoid is about 45 sec for cameras Af and An or Aa and An. A 6 m/sec error in the cross-track wind corresponds to a 1 pixel cross-track disparity error; in addition, the navigation accuracy from Level 1B2 is ±1 pixel. As with NM, the along-track strings are short enough to be considered aligned with the SOM grid, and only the orientation of the search window relative to the SOM grid must be taken into account.

The RS algorithm notes the pattern of relative brightness in a string by sorting the array of BRF's in the string and keeping track of the indices corresponding to their spatial locations. Comparison of the indices obtained from equally sized strings from two different view directions then determines whether there is a potential match. The algorithm is applied sequentially to different relative locations of one-dimensional strings of data within windows in the reference and comparison images. RS takes advantage of the fact that cloud-top heights frequently remain within ±250 m over distances of several kilometers to make the search efficient. A window (known as the filter window) is used in the reference image in addition to a search window in the comparison image in order to generate a number of potential matches, which are then median filtered. There are two reasons for this: (1) the raw results of RS are noisy, and (2) non-epipolarity, wind motion uncertainties, and image navigation uncertainties require allowing some variation in the location of the target string. Results from M2 and/or M3 at surrounding grid points, if available, are incorporated into the median filtering as well.

RS is based on matching the index tables (e.g., Numerical Recipes in C, section 8.3) of strings of BRF's from two different view zenith angles. These index tables essentially sort the BRF's from each of the strings in ascending order by their relative values; the indices correspond to the spatial locations of the BRF's. If two or more BRF's have identical values, they are assigned

the same rank in the index string. A match of the BRF patterns from the two different views means that they have identical index tables. Thus it is not necessary for, say, the mean value of the BRF's in the patches to be equal, just the relative brightness pattern. The index tables are constructed using simple inequality logic, and the test for a match is done by differencing the index tables.

When considering a particular string of pixels in the filter window, we attempt to match it with all along-track strings in the search window in the comparison image. A successful match of the target string occurs when (a) a single comparison string matches it, or (b) if multiple comparison strings match the target string. In the latter case, there is no metric to minimize for RS. Then, the multiple results are averaged, this average is returned as the match, and is input to the subsequent median filtering. Figure 13 depicts the manner in which RS is applied. Note that the target strings and comparison strings are rastered through their respective windows, beginning at a location in which the string center coincides with a pixel at a corner of the window. Thus, at times part of the string may extend beyond the window edge. Note also that the search window is parallel to the flight direction, and thus at a slight angle relative to the SOM grid coordinate system. However, the filter window and the target and comparison strings are aligned with the SOM grid.

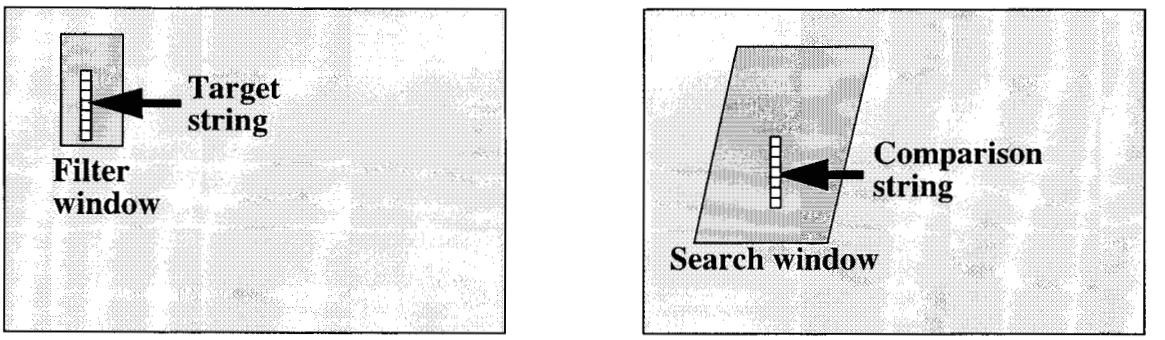

Figure 13. Terminology used in application of RS

In comparing an along-track string in the filter window with an along-track string in the search window, it is essential in applying RS that they be the same number of pixels in length. We require RS to have no tolerance for missing pixels. That is, if any string contains a pixel for which the RDQI > $RDQI_3$, that string is not considered in the matching process.

Each of the pixels in the along-track strings is provided with an index, numbering them sequentially from 1 to N. The assignment of indices is arbitrary, as long as the same assignment is made in the target and comparison strings. The N indices are then sorted as described above, and then the following quantity is calculated:

$$S_{RS} = \sum_{n=1}^{N} |I_{n,\,comp} - I_{n,\,targ}| \tag{18}$$

where $I_{n,comp}$ is the index of the n^{th}-most BRF in the comparison string, in which $n = 1$ corresponds to the darkest BRF value and $n = N$ corresponds to the brightest, and $I_{n,targ}$ is the equivalent value in the target string. A match between a target string and a comparison string is deemed successful if S_{RS} is identically zero.

As described in §3.4.5.2.4, RS is applied only to those locations within the filter window which cannot be "filled in" by results from surrounding grid points that have been successfully matched using M2 and/or M3. That is, if nearby M2 or M3 results are available as a result of marching through all of the data in the working strip, and the target patches associated with the M2/M3 matches overlap the filter window, the overlapped pixels in the filter window are assigned the M2/M3 results. RS is then applied to the remaining pixels. After values of S_{RS} are calculated by sliding the comparison string around the search window, and using all target strings within the filter window, a set of x and y disparities is established from all locations of the target string within the filter window which resulted in successful matches. The median values of x and y (obtained either from neighboring M2/M3 results or from RS) are then calculated. Thus, when applying the median filter to the RS results, we will take advantage of the more accurate retrievals from M2 and M3 if they are present within the filter range. Median filtering is relatively insensitive to the numerical values of the noise in a data set. Unlike other common smoothing techniques, median filters do not blur the locations of boundaries and thus are particularly useful in treating cloud fields without changing apparent cloud fraction. In applying the median filter to points at the swath edges, the left or right half of the filter window is missing. This is dealt with simply by calculating the median from among the disparities that are available.

3.4.1.4 Stereo matcher applicability test

The stereo matching schemes described above will fail if the areas to be matched are featureless, i.e., devoid of spatial contrasts. Thus, prior to establishing a valid stereo match, it is necessary to insure that sufficient contrast is present. Because the M2 matcher is used as part of both the cloud motion and height retrievals discussed below, we devise a stereo matcher applicability test that is invoked just prior to applying the M2 matcher. This test is performed on the target patch of the M2 matcher. To pass the stereo matcher applicability test, we require

$$q = \frac{\sigma_{M2} \cdot SNR_{1x1}[\rho_{\langle R\rangle}] \cdot (R_{max} - R_{min})}{N_{pts} \cdot \langle R\rangle} \geq q_{thresh} \tag{19}$$

where q_{thresh} is a configurable threshold and $SNR_{1x1}[\rho_{\langle R\rangle}]$ is the instrument signal-to-noise ratio for unaveraged (1x1) data, evaluated at the equivalent reflectance $\rho_{\langle R\rangle}$ that corresponds to the

mean BRF $\langle R \rangle$. The value of σ_{M2} is given by Eq. (11). We note that $\rho_{\langle R \rangle} = \mu_0 \langle R \rangle$, so we obtain from the Ancillary Radiometric Product the table of signal-to-noise ratios, SNR_{am}, for the tabulated set of equivalent reflectance values for the appropriate camera and the red band and averaging mode am = 1x1, and linearly interpolate these to the required value of $\rho_{\langle R \rangle}$. The expression of Eq. (19) implies that if the variation in BRF within the target patch is solely due to instrument radiometric random noise (i.e., $R(x_i, y_j) \cong \langle R \rangle + noise$) and not actual scene contrast, $q \sim 1$. We require the contrast to be at least a factor q_{thresh} larger than the noise. We set $q_{thresh} = 5$.

When applying Eq. (19) during wind retrieval, and the test passes, the local maximum around which the target patch is constructed is deemed usable for matching. If the test fails, the local maximum around which the target patch is constructed is deemed unusable for matching. This condition is applied during the stereo-matching for both wind and height retrieval. If a target patch fails the contrast test during application of M2 within NM, the matching process is aborted for that target local maximum and no disparity is assigned. When applying the test during height retrieval, and the test passes, application of M2, M3, and RS is permitted to proceed. If the test fails, the subregion containing the center of the target patch is labelled as having no disparity available, and no attempt to apply M2, M3, or RS is made.

3.4.2 Select cameras for stereo processing

3.4.2.1 Physics of the problem

As discussed in §3.4.1, triplets of cameras are required for stereoscopic wind retrieval. Each triplet must be asymmetrically positioned in view zenith angle, and a wide spread in view zenith angle is desirable. Additionally, it is desirable to use two sets of triplets in order to obtain more confidence in the retrievals. Thus, the triplets An-Bf-Df and An-Ba-Da are most suitable. For stereo height retrievals, two camera pairs are used. To maximize the coverage and obtain a continuous height field (as opposed to wind retrieval, where sparse coverage is adequate), we desire cameras which are expected to exhibit the greatest degree of correlation and the sharpest imagery. Thus, the An-Af and An-Aa pairs are best suited to this.

3.4.2.2 Mathematical description of the algorithm

For several reasons, it is possible that the ideal set of cameras is not available. Thus, it is necessary to build flexibility into the data processing approach such that the software will be able to deal with such contingencies. This is done by allowing the reference and comparison cameras to be modified through the use of the Configuration File (Table 9), and building the processing software in a generalized way to enable a different choice of cameras.

3.4.3 Perform stereo matching for cloud motion retrieval

3.4.3.1 Physics of the problem

A number of well distributed matching features over a domain where cloud motion is regarded as constant are required to provide an accurate and reliable cloud motion retrieval. A strategy to combine the merit of both the feature-based and area-based matching techniques is employed. The matching algorithm is aimed at identifying accurately a small subset of cloud features from the reference, comparison, and common comparison camera triplets. This matching algorithm has the ability to search for conjugate features in two dimensions. This is required because the existence of unknown cloud advection can cause disparities in both the along-track and cross-track directions. First, NM is applied to BRF imagery in the red channel at the original 275 m resolution for each along-track column. Following application of NM, we use M2 to resolve any ambiguities in the matches as described in §3.4.1.3.1.

3.4.3.2 Mathematical description of the algorithm

Cloud motion is retrieved using stereo-matchable features whose physical location, at the time of observation by the wind retrieval common comparison camera, places them within a volume of orbital swath data encompassing a strip that is 70.4 km long in the along-track direction, and which contains the width of the swath in the cross-track direction. The feature must also lie within a height range that extends from the surface ellipsoid to a level 20 km above the surface ellipsoid. The strips for the camera triplets are also taken to be the width of the swath, and overlay the 70.4-km retrieval strip, but have increased dimensions in the along-track direction. This is necessary such that the search windows are large enough to encompass features which are displaced due to motion from one image to the other or due to a shift resulting from a height different than the surface ellipsoid. Specifically, the additional length accounts for the apparent cloud displacements associated with along-track disparities resulting from a maximum cloud height of 20 km above the surface ellipsoid, the actual displacements resulting from a maximum along-track wind speed of 100 m/sec, and an additional 4 pixels owing to image navigation uncertainties.

Figure 14 shows the example of the nominal case in which the reference camera is Bf (or Ba), the common comparison camera is An, and the comparison camera is Df (or Da). The leading and trailing areas for the off-nadir cameras are non-symmetrical with respect to the wind retrieval strip because the along-track disparity due to height displaces a feature in the direction away from the nadir point, whereas displacements due to along-track wind may be either toward or away from the nadir point. The working strip for An, despite its orthonormal view, is larger than the wind retrieval strip because it must include the area within which a feature viewed by the reference camera (Bf or Ba) may be found.

Taking the swath width to be W pixels, we require W x 275 m to be an integer multiple of

70.4 km, such that wind fields can be determined within mesoscale domains measuring $(70.4\ \text{km})^2$. The NM/M2 algorithm is now applied to the selected wind retrieval cameras within each domain to identify local maxima in the reference camera and candidate common comparison camera matches that are associated with each, and to identify local maxima in the reference camera and candidate comparison camera matches that are associated with each of those. As stated in §3.4.1.3.1, NM is used for all levels of local maxima down to the second lowest. In applying M2 here, a M_c pixel (cross-track) x M_a pixel (along-track) patch is constructed around each local maximum in the reference image, and M_c x M_a patches are constructed around each candidate local maximum in the comparison image, in order to find out if a unique match can be found. We establish nominal values of $M_c = 11$ and $M_a = 21$. The values of M_c and M_a are chosen to be odd numbers such that the pixels at which the local maxima occur are located at the center of their respective patches. If multiple matches satisfy the M2 thresholding, the ambiguity test described in §3.4.1.2.1 is applied to establish whether a viable match exists. We note that all local maxima are marked "unused" prior to each invocation of the NM algorithm.

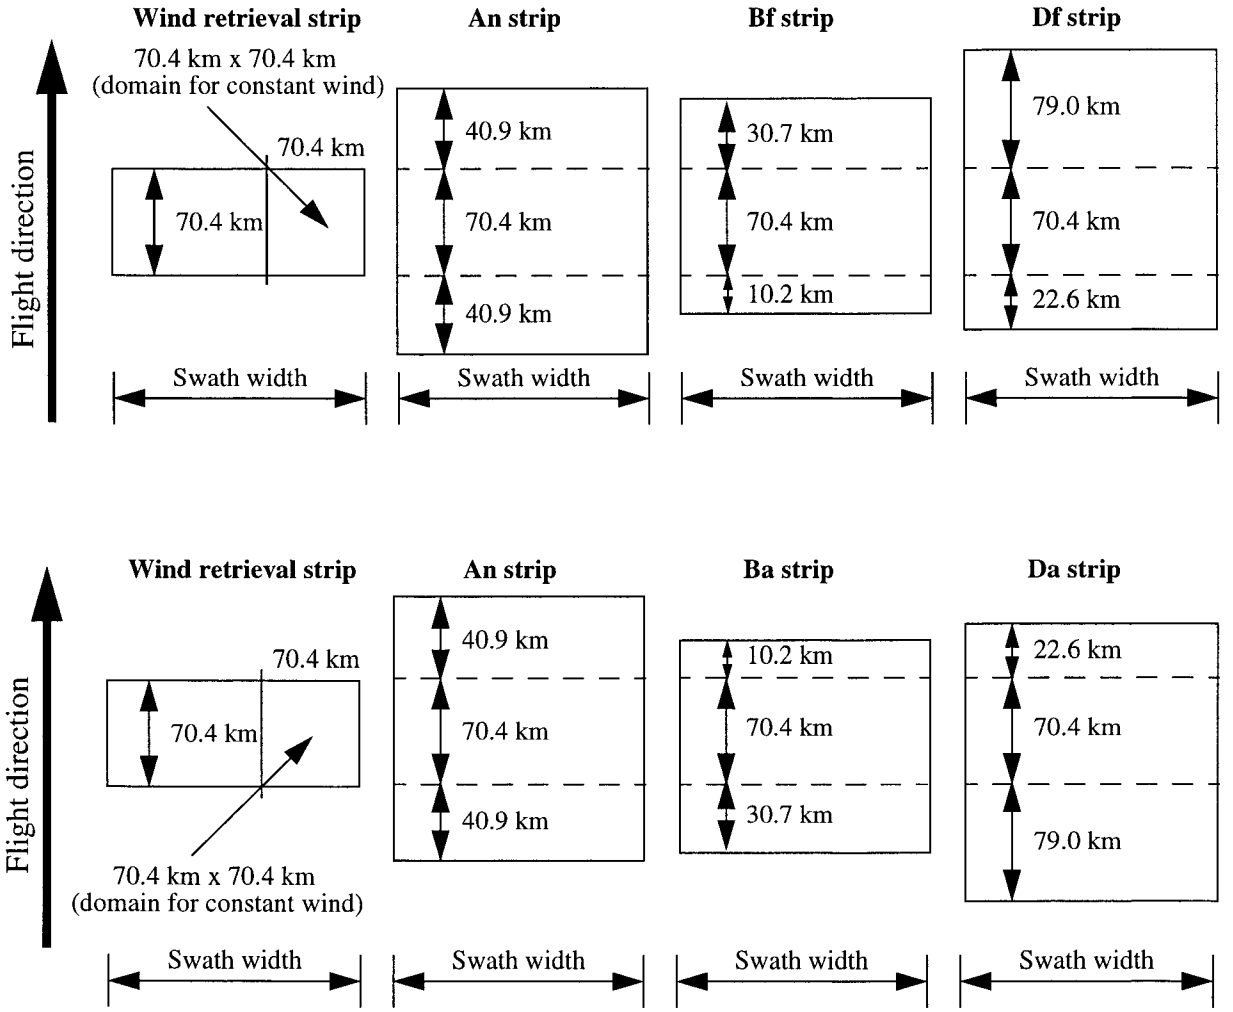

Figure 14. Dimensions of working strips for the An, B, and D cameras for wind retrieval

Next, the points that have both a reference-common comparison (e.g., Bf-An) and reference-comparison (e.g, Bf-Df) match associated with them are retained. All others are discarded. This process is performed for both the forward and aftward triplets. Finally, all remaining matches within a given $(70.4\text{-km})^2$ domain are used to determine cloud motion using a three-dimensional ray intersection algorithm described in the next section.

3.4.4 Retrieve cloud motion

3.4.4.1 Physics of the problem

Clouds at different altitudes may move at different velocities, and the general solution may require a continuum of altitudes and velocities. Over mesoscale domains, however, we can minimize this complexity by seeking only the two most common cloud velocities, finding the corresponding median altitudes associated with these two velocities, and then taking an average of these two heights to establish an altitude designated h_{wind} which separates the atmosphere into two altitude ranges corresponding to the two most common cloud displacement vectors. Note that it is possible, e.g., in the case of partially cloud-free land, for one of the two most common velocities to be a value close to zero, owing to matches obtained from the surface. In this case the two altitude bins will correspond to a layer comprised of the surface terrain and a layer containing cloud. In other situations, the two altitude bins may correspond to two distinct cloud layers.

3.4.4.2 Mathematical description of the algorithm

The first step is to use the disparities determined from the NM/M2 stereo matchers to calculate along-track wind speed, cross-track wind speed, and a height associated with each. Assume that there is no vertical cloud motion but only a constant horizontal cloud motion. Then, the cloud motion vector, the conjugate look rays from image matching of the common comparison-reference-comparison camera triplets, and the surface disparities resulting from the matching should form a closed loop in a three-dimensional Cartesian coordinate system as shown in Figure 15.

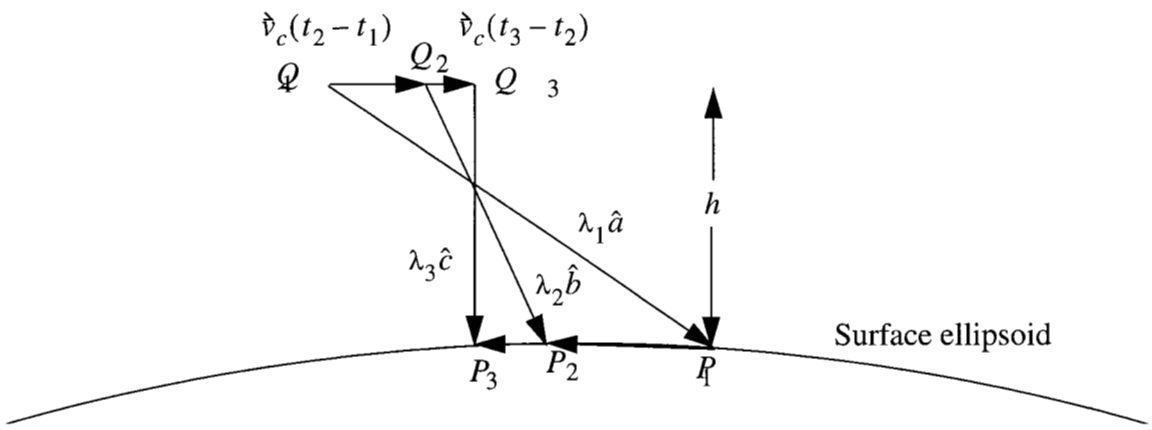

Figure 15. Ray intersection to retrieve cloud motion

To increase the accuracy of the wind calculations, all of the incoming values of latitude, longitude, look vectors and timetags must be linearly interpolated to the center of the 275-m pixel location for each camera, rather than simply supplied at the nearest 1.1-km or 2.2-km grid center. When necessary, extrapolation rather than interpolation is used.

The following equations represent this relationship:

$$\overrightarrow{P_1P_3} = \lambda_3\hat{c} + \vec{v}_c(t_3 - t_1) - \lambda_1\hat{a} \tag{20}$$

$$\overrightarrow{P_2P_3} = \lambda_3\hat{c} + \vec{v}_c(t_3 - t_2) - \lambda_2\hat{b} \tag{21}$$

where Q_1, Q_2, and Q_3 are the locations of a cloud feature at times t_1, t_2, and t_3, respectively, and $P_1(x_1,y_1,z_1)$, $P_2(x_2,y_2,z_2)$, and $P_3(x_3,y_3,z_3)$ are the conjugate image locations of the same cloud feature on the surface ellipsoid in the Cartesian coordinate system; $\hat{a}$, $\hat{b}$, and $\hat{c}$ are the unit vectors of the conjugate look rays; and λ_1, λ_2, and λ_3 are the unknown scale factors of the look rays for them to intersect with the cloud, respectively. For each set of 3 conjugate look rays, the six unknowns describing the cloud motion by the velocity $\vec{v}_c$ and the cloud locations by the ray vector scale factors λ_1, λ_2, and λ_3 can be solved simultaneously with these six equations. However, in order to easily impose the condition $v_{cz} = 0$ into Eqs. (20) and (21), it is preferable to represent this relationship in a local coordinate system where its z-axis is aligned with the zenith direction at the common comparison camera image point P_3, and having its origin at P_3. The x-axis and the y-axis can be defined as in the along-track and the cross-track directions, respectively. Because of the condition $v_{cz} = 0$,

$$\begin{cases} \lambda_1 = \dfrac{c_z}{a_z}\lambda_3 - \dfrac{(z_3 - z_1)}{a_z} \\ \lambda_2 = \dfrac{c_z}{b_z}\lambda_3 - \dfrac{(z_3 - z_2)}{b_z} \end{cases} \tag{22}$$

Equations (20) and (21) can be decomposed into a set for along-track and another set for cross-track:

$$\begin{cases} (x_3 - x_1) - \dfrac{a_x}{a_z}(z_3 - z_1) = v_{cx}(t_3 - t_1) + \lambda_3\left(c_x - c_z\dfrac{a_x}{a_z}\right) \\ (x_3 - x_2) - \dfrac{b_x}{b_z}(z_3 - z_2) = v_{cx}(t_3 - t_2) + \lambda_3\left(c_x - c_z\dfrac{b_x}{b_z}\right) \end{cases} \tag{23}$$

$$\begin{cases} (y_3 - y_1) - \dfrac{a_y}{a_z}(z_3 - z_1) = v_{cy}(t_3 - t_1) + \lambda_3\left(c_y - c_z\dfrac{a_y}{a_z}\right) \\ (y_3 - y_2) - \dfrac{b_y}{b_z}(z_3 - z_2) = v_{cy}(t_3 - t_2) + \lambda_3\left(c_y - c_z\dfrac{b_y}{b_z}\right) \end{cases} \tag{24}$$

Equations (23) and (24) are equivalent to Eq. (6), but generalized to three dimensions. Because we have four equations and three unknowns (λ_3, v_{cx}, and v_{cy}), the problem is overdetermined and the equations are solved by least-squares.

The solutions of Eqs. (23) and (24) provide directly the cloud velocity $\vec{v}_c$ and its locations defined by λ_1, λ_2, and λ_3 at the captured times. The height of the cloud of this matched triplet is obtained by projecting the common comparison camera view look ray, $\lambda_3\hat{c}$, onto the normal direction at the surface ellipsoid where the look ray intersects:

$$h = \lambda_3 c_z \tag{25}$$

Whenever h returns a value < 0 km or > 20 km, or whenever the horizontal location of the feature at the time of observation by the common comparison camera places it outside the $(70.4\text{ km})^2$ domain of interest, the associated points are not used for further processing.

For all successful matches, the components of the cloud velocity vectors $\vec{v}_c$ and the associated heights are now used to determine the two most common cloud velocities within two altitude ranges separated by the height h_{wind}. The procedure for determining these two velocities and the value of h_{wind} is as follows. The independent results from the forward triplet matches and from the aftward triplet matches, if available, are used together.

(1) Sort the cloud velocity vectors into a two-dimensional histogram, where each bin of the histogram corresponds to cross-track and along-track wind speeds. The bin widths are 6 m/sec in both directions. In each direction the bins have speeds ranging from - (96 - 102) m/sec up to + (96 - 102) m/sec, with an additional bin for speeds with absolute values > 102 m/sec. The histogram therefore contains 36 x 36 bins.

(2) Count the number of points in each bin. Let n_1 be the highest count found, and let n_2 be the next highest count found.

(3) Determine how many bins have count n_1.

  (a) If only one bin has this count, establish this as the most populated bin. Go to step (4).

  (b) If exactly two bins have this count, establish these as the two most populated bins. Go to step (5).

(c) If more than two bins have this count, concatenate together all bins having count n_1 for which the cross-track wind speed is ≥ 0, and separately concatenate together all bins having count n_1 for which the cross-track wind speed is < 0. Depending on whether the cross-track wind speeds are all ≥ 0, all < 0, or both, this will lead to either one or two concatenated bins.

(i) If this results in a single concatenated bin, establish this concatenated bin, and the set of all heights and velocities within it, as the most populated bin. Go to step (4).

(ii) If this results in two concatenated bins, establish these two concatenated bins, and the set of all heights and velocities within them, as the two most populated bins. Go to step (5).

(4) Determine how many bins have count n_2.

(a) If only one bin has this count, establish this as the second most populated bin.

(b) If two or more bins have this count, concatenate together all bins having count n_2 for which the cross-track wind speed is ≥ 0, and separately concatenate together all bins having count n_2 for which the cross-track wind speed is < 0.

(i) If this results in a single concatenated bin, establish this concatenated bin, and the set of all heights and velocities within it, as the second most populated bin.

(ii) If this results in two concatenated bins, establish as the second most populated bin the concatenated bin that had the larger number of individual bins contributing to it. Include all heights and velocities from the individual bins in the concatenated bin.

(iii) If both concatenated bins had the same number of individual bins contributing to them, establish as the second most populated bin the set of heights and velocities from all individual bins having count n_2.

(5) Average the velocities within each of the two resulting "most populated" bins. Call these averages $\langle \vec{v_c} \rangle_1$ and $\langle \vec{v_c} \rangle_2$.

(6) Calculate h_{wind} from

$$h_{wind} = \frac{median\{h_1\} + median\{h_2\}}{2} \tag{26}$$

where $\{h_1\}$ and $\{h_2\}$ are the set of heights in each of the two bins.

Now, if $median\{h_1\} \geq h_{wind}$, then bin 1 is designated the High Cloud bin, $\langle \vec{v_c} \rangle_1$ is assigned to this bin, and $\langle \vec{v_c} \rangle_2$ is assigned to the Low Cloud bin. Alternatively, if $median\{h_1\} < h_{wind}$, then bin 1 is designated the Low Cloud bin, $\langle \vec{v_c} \rangle_1$ is assigned to this bin, and $\langle \vec{v_c} \rangle_2$ is assigned to the High Cloud bin. The height range of the Low Cloud bin is established as $0 \leq h < h_{wind}$, and the height

range of the High Cloud bin is $h_{wind} \leq h \leq 20$ km.

In the event that all of the wind velocities fall within a single bin, such that there is no bin with count n_2, we set $\langle \vec{v_c} \rangle_2 = \langle \vec{v_c} \rangle_1$, the average wind velocity, and set $h_{wind} = 10$ km.

If the total number of points from both the forward and aftward wind disparity retrievals is less than a configurable value (currently set to 100), the above procedure is replaced with a default. In this case, we set $h_{wind} = 10$ km, and the wind velocity in both the Low Cloud and High Cloud bins is set to 0.

3.4.4.3 Example of cloud motion retrieval

In order to retrieve cloud-track winds, cloudy scenes were generated in the following manner. First, a cloud height field of 256 x 256 pixels (70.4 km) was created with a fractal cloud generator. The cloud field was assumed to consists of 256 x 256 prisms. Each one of these prisms had a base area of 275 x 275 m^2 and a height drawn from the previously generated fractal cloud height field. Then an arbitrary brightness field was assigned to the lambertian cloud tops with higher tops being brighter. This cloud field with all the prisms being at the same cloud base height was assumed to move with a constant horizontal velocity. Finally, a brightness value was assigned to each pixel of a 256 x 256-pixel region on the surface ellipsoid by means of a ray tracing algorithm. The input to the ray tracing consisted of the geodetic latitude and longitude, and the corresponding camera look-vectors and timetags of each pixel. In the course of the ray tracing, the brightness value was linearly interpolated from neighboring values whenever the side of a cloud column was hit instead of its top.

An example of the cloud motion retrieval is plotted in Figure 16. The results given below were obtained for 100% cloud cover. The cloud base height was 1 km with cloud top heights varying between 1 km and 3.8 km. The clouds were moving in the NE direction, thus the N-S and E-W velocity components were identical and varied from 0 to 50 m/s with 1 m/s increments. Here the solid lines represent the true cloud speeds while the plus signs mark the retrieved values. The accuracy is ±3 - 4 m/sec in the N-S (approximately the along-track) direction. The retrieved E-W (approximately the cross-track) components follow the true values very closely with a typical error of 1 - 2 m/s. The root-mean-square error corresponding to the retrievals depicted in Figure 16. Retrieval Results for a single cloud layer was $rmse_{NM} = 2.2$ m/sec.

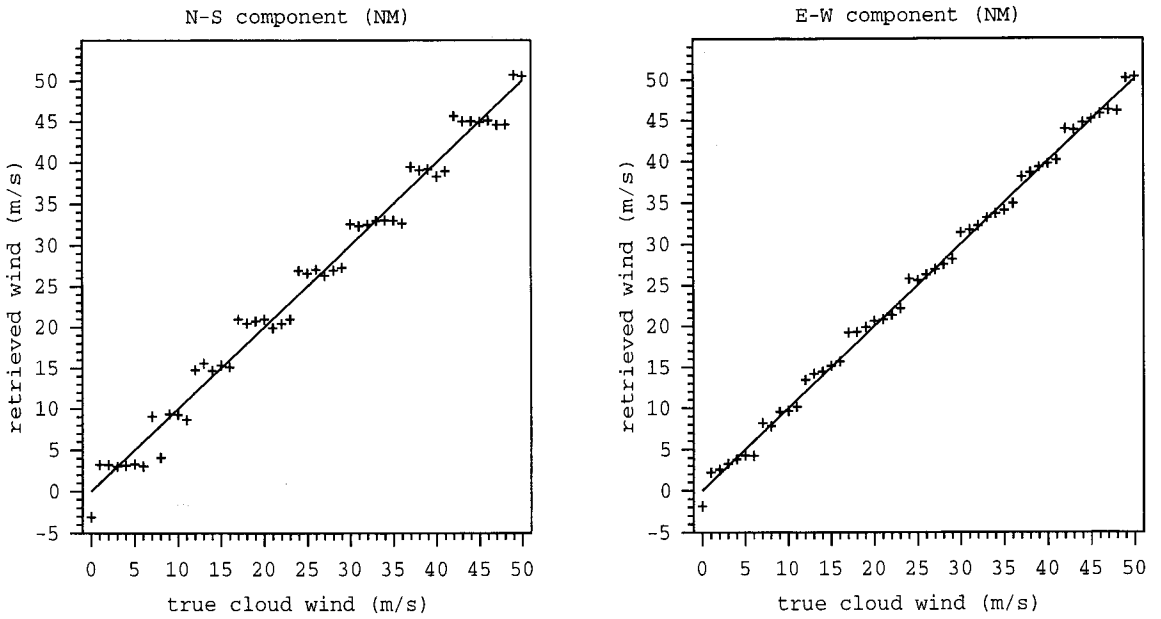

Figure 16. Retrieval results for a single cloud layer

3.4.5 Perform stereo matching for height retrieval

3.4.5.1 Physics of the problem

The cloud velocity as determined above may be assumed constant over a domain of mesocale size. However, h is likely to be much more variable, and the total disparity must be retrieved at higher resolution in order to produce the albedos and BRF's at a spatial resolution of $(2.2\text{km})^2$. The assumption that there is a systematic and therefore removable effect due to the mesoscale wind components is not, however, critical to the ultimate determination of total disparities at fine resolution. It does affect how total disparities are interpreted (height versus wind displacement), and by removing the background disparity due to wind displacement before seeking smaller scale departures from the background shift, it reduces the search window, giving more consistently accurate results with fewer computations.

An approach combining both feature-based and area-based matching techniques is used, employing data from the height retrieval camera pairs (ideally, An-Af and An-Aa, because these views provide the highest resolution and greatest likelihood of identifying the same features from their relative brightness patterns). The input data used to obtain the disparities at high resolution are the unaveraged band 3 (red) BRF's referenced to the surface ellipsoid.

The approach uses a combination of the area-based matchers, M2 and M3, supplemented by the feature-based matching technique, RS, to fill in the coverage. These matchers are applied in a

hierarchical fashion. Coverage increases in going from M2 to M3 to RS, but accuracy decreases in the same order. M2 is used preferentially, then M3, and finally RS.

3.4.5.2 Mathematical description of the algorithm

In deriving height fields that are going to be used to determine values of RLRA, the imagery from the height retrieval reference camera is used as a reference. We establish an area in the reference image measuring 70.4 km or 256 lines in the along-track direction, and extending across the entire swath in the cross-track direction. As with the wind retrievals, the width of the swath is taken to be W pixels. This strip is divided into (70.4-km)^2 domains (each having a wind retrieval associated with it), and these domains each contain a 64 x 64 array of (1.1 km)^2 subregions. In order to allow for displacement (due to wind) and size (to solve for height) of the search windows, we provide additional length in the along-track direction for the imagery from the height retrieval comparison cameras on both the leading and trailing sides, using the same assumptions as in the first paragraph of §3.4.3.2. Figure 17 is an example for the case in which An is the reference camera and Af, Aa are the comparison cameras, for all domains across the strip.

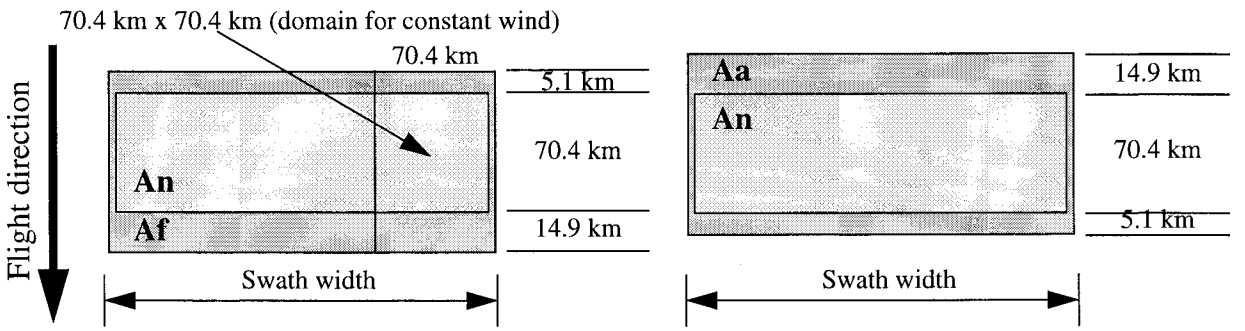

Figure 17. Dimensions of working strips for the An, Af, and Aa cameras for height retrieval

Because the height (and consequently, the disparity) field is expected to be correlated over small spatial scales, the efficiency of the stereo matching is improved by using previously determined neighboring matches, if available, as starting points for searching for a match. If either an adjacent along-track or cross-track disparity, or both, is available, the along- and cross-track disparities are used as the center points of small search windows. The union of these search windows is then used as the area to search for a best match, first using M2 and then, if required, M3. The procedure is referred to as the "previous match" method. If this process fails to establish a good match, a more extensive search window is constructed, as follows.

For each subregion in the reference image, the cloud velocities in the two altitude bins at this location are used to predict the location of search windows in the comparison image (initially the one forward of the reference, and then, independently, the one aftward of the reference), and the height ranges of the two bins (0 - h_{wind} and h_{wind} - 20 km) are used to establish the sizes of the search windows. If the absolute value of the difference in along-track components of $\langle \vec{v}_c \rangle_1$ and

$\langle \vec{v_c} \rangle_2$ is ≤ 12 m/sec, and the absolute value of the difference in cross-track components of $\langle \vec{v_c} \rangle_1$ and $\langle \vec{v_c} \rangle_2$ is also ≤ 12 m/sec, only the average velocity needs to be used to establish the location of the search window for the height retrievals. This is done because the disparity caused by a velocity difference of 12 m/sec is less than 2 pixels. In any case, we perform image matching over the area established by the union of the two search windows. The resulting disparity values are later converted into cloud heights. If the "best" disparity is found within an area of one of the search windows that does not overlap the other search window, the value of wind velocity associated with that window is used in the disparity-to-height conversion. If the best disparity is within an overlap region of the two windows, the average velocity of the two bins will be used.

Once the search windows associated with each subregion have been located, we establish whether sufficient data exist within it to proceed with application of the matchers. This is to avoid biases in the final result, that is, to insure that there is adequate opportunity to locate the correct match. To do this, we let N_{search} be the number of red band 275-m pixels within the search window, regardless of whether there is valid data in each of them. Let $N_{search,comp}$ be the number of red band 275-m pixels for the comparison camera to which the matchers are being applied that actually have valid BRF values. We require

$$\frac{N_{search,\,comp}}{N_{search}} \geq f_{search} \tag{26a}$$

where f_{search} is a minimum required amount of available data. This is a configurable parameter, and is set to 0.75.

If the requirement of Eq. (26a) is satisfied, the M2 matcher is applied. For those subregions for which M2 does not return a successful match, the M3 algorithm is applied. The previous match method is attempted first, with the image pyramid method then being attempted if that fails. All $(70.4\text{-km})^2$ domains across the working strip are processed in this fashion, to establish a set of disparities on 1.1-km spacing for the entire strip. For those points on this 1.1-km grid where neither M2 nor M3 established a match, we construct a filter window around the point, incorporate any available M2/M3 disparities from neighboring grid points, and subject the remaining points within the filter window to the RS algorithm. The set of disparities are then median filtered. If this procedure still fails to find a match, the grid point is flagged as "no disparity available". Following this procedure, each domain will have a 64 x 64 array of x and y displacements (disparities) from each camera pair. These are then converted to heights with confidence designators using the algorithm described in §3.4.6.

Note that when M2/M3 disparities from neighboring grid points are used to "fill in" the filter window, it is permissible to use M2/M3 values from an adjacent $(70.4\text{-km})^2$ domain, but only if the disparities are derived using the same camera pair.

3.4.5.2.1 Establish array of target patches and associated search windows

Within a domain in the height retrieval reference camera image, a 64 x 64 array of subregions measuring $(1.1\ \text{km})^2$ is identified. The centers of each subregion are separated by 1.1 km in both the cross-track and along-track direction. Centered on each subregion we place a N_c pixel (cross-track) x N_a pixel (along-track) target patch. We establish nominal values of $N_c = 6$ and $N_a = 10$. Even values are chosen such that the patch is centered on an SOM grid intersection. Matching to each patch is performed by establishing a search window in either the comparison camera imagery, with matches performed on the reference-forward comparison pair and the reference-aftward comparison pair independently.

The locations of the search windows based on previously established matches are established as follows. If a previously evaluated neighboring subregion in the along-track direction has a valid match, a fiducial point is located at a point in the comparison image that is offset by the along-track and cross-track disparities of that match from the SOM coordinates of the geographic center of the current target patch. A small search window of dimensions B_cxB_a is established around this point. Similarly, if a previously evaluated neighboring subregion in the cross-track direction has a valid match, a fiducial point is located at a point in the comparison image that is offset by its along-track and cross-track disparities from the SOM coordinates of the geographic center of the current target patch. A small search window of dimensions B_cxB_a is also established around this point. The union of these two small search windows then comprises the area within which to apply the area-based matchers. We nominally set $B_c = 5$ pixels and $B_a = 7$ pixels.

If the neighboring along-track and cross-track subregions have previously determined matches associated with them, these matches are considered to be "valid" for the above purpose if the value of the matcher metric satisfies the following:

$$S_{M2}(\text{previous match}) \le f_2 \cdot T_{M2} \quad \text{if the M2 matcher was used for the previous match}$$

$$S_{M3}(\text{previous match}) \le f_3 \cdot T_{M3} \quad \text{if the M3 matcher was used for the previous match} \qquad (26b)$$

By setting f_2 and f_3 to values < 1.0, this criterion places a stringent requirement on the previously established matches. We establish nominal values of $f_2 = 0.5$ and $f_3 = 0.5$.

The M2 matcher is then applied to the search window resulting from this process, and if multiple matches are found, the uniqueness test described in 3.4.1.2.1 is used. If a unique match is not found, the M3 matcher is tried. If this method fails to identify a match, we default to using a more extensive search window, coupled with a pyramid scheme for finding the match. As before, we first attempt M2 and then proceed to M3 if necessary.

The locations of the large search windows are established as follows. We first locate a fiducial point within the search window. Starting with the location in the comparison image that has the same SOM coordinates as the geographic center of the target patch, the fiducial point is offset from this grid intersection by an along-track (x-direction) and cross-track (y-direction) number of pixels determined by the components of the wind vector:

$$\Delta x_f = -\langle \vec{v_c} \rangle_x \Delta t_f \qquad \Delta y_f = -\langle \vec{v_c} \rangle_y \Delta t_f \tag{27}$$

$$\Delta x_a = \langle \vec{v_c} \rangle_x \Delta t_a \qquad \Delta y_a = \langle \vec{v_c} \rangle_y \Delta t_a \tag{28}$$

where Δt_f is the absolute value of the time difference between viewing a particular feature in the reference and forward comparison cameras; Δt_a is similarly defined for the reference and aftward comparison cameras. These are obtained by taking the Δt values shown in Table 4, expressed as line numbers, and multiplying by 40.8 msec, the MISR camera line repeat time. For example, Δt is approximately 45.4 sec between the nadir and off-nadir A cameras. Note that the displacement is negative in Eq. (27) because the feature is seen earlier in the camera that views more forward than the reference camera, and positive in Eq. (28) because it is seen later in the more aftward camera.

We now construct a search window around the fiducial point. The dimensions of the search window are determined from the limits on the height bins of the wind vectors, as shown in Figure 18. The distance l is established to take into account non-epipolarity, uncertainties in image registration, and uncertainties in the wind retrieval. We set l = 1.1 km, and the width of the window in the cross-track direction is equal to $2l$. The along-track distance d, which leads the fiducial point in the forward comparison image and trails it in the aftward comparison image, is given by

$$d = l + h_2 \left| \tan\theta_c - \tan\theta_r \right| \tag{29}$$

where h_2 is the upper bound altitude of the wind velocity bin ($h_2 = h_{wind}$ for the Low Cloud bin and h_2 = 20 km for the High Cloud bin), θ_c is the nominal view zenith angle of the comparison camera relative to the local normal, and θ_r is the nominal view zenith angle of the reference camera relative to the local normal. The along-track length of the search window, D, is given by

$$D = 2l + (h_2 - h_1) \left| \tan\theta_c - \tan\theta_r \right| \tag{30}$$

where h_1 is the lower bound altitude of the wind velocity bin (h_1 = 0 for the Low Cloud bin and h_1 = h_{wind} for the High Cloud bin). In order to obtain the correct results from Eqs. (29) and (30), the θ angles are specified to be positive for camera look vectors viewing forward of nadir, and negative for camera look vectors viewing aftward of nadir, as shown in Table 4. Note that the search window is longer in the along-track direction due to the greater disparities which result from the

camera geometry, and we include the additional length $2l$ to account for uncertainties, as described above. Searching within this window is performed on the comparison cameras to establish residual disparities resulting from heights differing from these reference levels. It is possible for the length of the search window, D, to be short enough such that both ends of the window are on the same side of the fiducial point in the along-track direction.

Conversion of the dimensions l, d, and D to pixels is accomplished by dividing by 275 m, and rounding up to the nearest integer. Thus, the search window measures L_c pixels cross-track x L_a pixels along-track, where $L_c = 2l/(275\text{ m})$ and $L_a = D/(275\text{ m})$, rounded up to the nearest integer. For points on the working strip where M2 and M3 do not return results, the locations and dimensions of the search windows are temporarily stored such that when RS is applied it is not necessary to recalculate them.

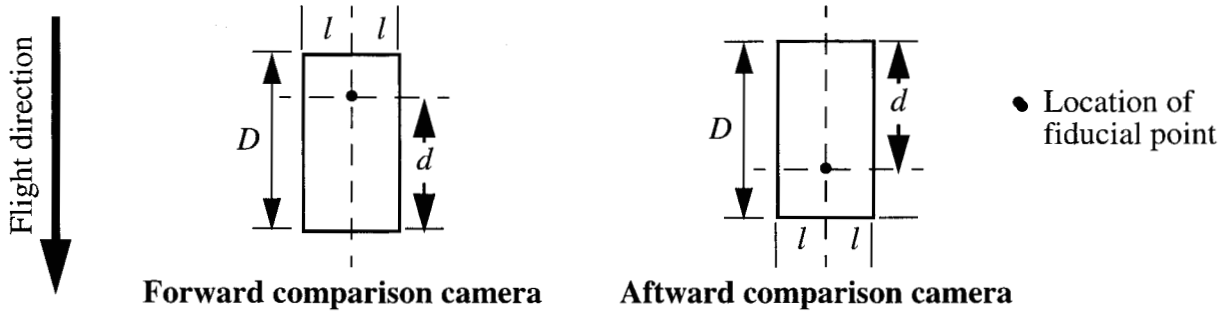

Figure 18. Dimensions of search windows for finding height disparities

As a final step in setting up the search window, it is rotated to the appropriate orientation within the SOM block.

3.4.5.2.2 Apply M2 algorithm

Superimposed on the L_c x L_a search window, individual N_c x N_a arrays of 275-m pixels are identified. These are the comparison patches. (Recall that the search window is tilted within the SOM coordinate system, but the patches are aligned with the SOM grid. The search window is also scaled to the appropriate size based on the pyramid level.) The comparison patches overlap, and are displaced from one another by 275 m in each of the cross-track and along-track directions. The comparison patches are successively centered on each of the grid intersections contained within the search window, including those on the window edges. Thus, at the last level of the pyramid, there are approximately $(L_c + 1)$ x $(L_a + 1)$ such patches within the search window, where the approximate nature arises due to roundoff caused by the tilted search window. At the coarser level, the count decreases by a factor of 4. Since we match the centers of each patch, the patches themselves will extend beyond the dimensions of the search window at the edges of the window. Each comparison patch is compared to the target patch from the reference camera using M2. We test whether a value of $S_{M2} \leq T_{M2}$ exists to determine if a successful match is achieved. If multiple matches that

pass the threshold are obtained, we resolve the ambiguity by applying the test described in §3.4.1.2.1. Based on experiments with test data we set $T_{M2} = 0.75$. This value may need to be adjusted once actual MISR data become available.

If a unique match is found at the initial level of the pyramid, a new search window of dimensions L_c' x L_a' is centered on the location of the match and used for searching at the next level of the pyramid. We nominally set $L_c' = 5$ pixels and $L_a' = 5$ pixels. The ambiguity test is applied at all but the highest resolution levels of the pyramid.

If a successful match (the one with the smallest matcher metric) is found at the final level of the pyramid, we move to the next subregion (and embedded target patch) in the reference image. Otherwise, the procedure is repeated using M3.

At the coarsest level of the pyramid, the M2 algorithm is carried out on the area which is the union of the two search windows corresponding to the wind velocities in the Low Cloud and High Cloud bins. First, we establish and march through the search window for the Low Cloud height bin and note the disparities and associated metric values for all points that meet the threshold cut. The same procedure is then repeated for that portion of the window associated with the High Cloud height bin that is distinct from the window associated with the Low Cloud bin, and the list of possible matches is augmented. The best match is then selected from the combined list, applying the ambiguity test if necessary. If the wind speeds in the two bins are sufficiently close as described in §3.4.5.2, the algorithm runs once as described within a single search window defined by the average velocity and this averaged velocity is used in the subsequent retrieval of height feature from disparities.

3.4.5.2.3 Apply M3 algorithm

This step is conditional upon the results of §3.4.5.2.2. If no match was found by M2 at a given level of the pyramid, the procedure described in §3.4.5.2.2 is repeated, this time using M3. We set $T_{M3} = 1.0$; again this value may be revised after launch. At the next finer level of the pyramid, we again revert to using M2, followed by M3 if necessary, basing the search window location on the match location from the previous level of the pyramid, regardless of whether it was provided by M2 or M3. If a successful match is found at the finest level of the pyramid, we move to the next target patch in the reference image. However, if the final match (by whichever method) was determined by M2, a verification check is performed on this single point to determine whether M3 considers this to be a good match as well. If so, this match will be given higher weight in subsequent processing. The combination of M2 and M3 is applied to each grid location in the working strip. This process is completed before consideration is given to using RS.

3.4.5.2.4 Apply RS algorithm

After M2 and M3 are applied to the entire working strip, each 1.1-km subregion in the G_c x G_a array of disparities which does not have a disparity associated with it is re-examined, this time using the RS algorithm. It is necessary to wait until the grid is filled out before applying this step because any neighboring successful M2/M3 retrievals will be combined with the RS results. To apply RS, we establish a filter window of dimension of J_c x J_a pixels, centered on the 1.1-km sub-region under consideration. This window is depicted in Figure 19 and is aligned with the SOM grid. In order to keep the filter window centered on the same grid intersection as the target patch, we choose J_c and J_a to be even numbers. Nominally, we set $J_c = J_a = 6$.

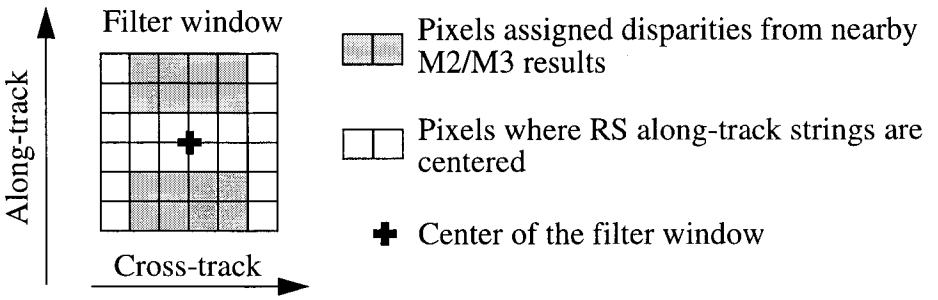

Figure 19. Method of establishing M2/M3/RS disparities within the filter window

The following steps are now performed:

(1) Establish pixels within the filter window that overlap any of the N_c x N_a target patches of surrounding grid points that have successful M2 or M3 matches from the same pair of cameras to which RS is being applied. Assign the M2/M3 disparities to these locations. If any of these pixels are overlapped by more than one M2/M3 target patch, the disparity values assigned are determined by:

(a) Giving M2 verified by M3 precedence over M2, and M2 preference over M3;

(b) Using the smallest value of the matcher metric to choose the preferred value. This means that the M2/M3 matcher metrics will need to have been stored in a temporary array.

(2) For all remaining pixels in the filter window, center a K_a-pixel long target string on each of those locations and apply RS. The comparison strings are similarly located within the search window(s). The value of K_a is an odd number, such that the strings are centered on pixels. Nominally, $K_a = 7$.

(3) Out of the J_c x J_a pixels in Figure 19, determine whether at least a pre-established proportion of them (a configurable value currently set to 50%) now have disparity values assigned to them. If so, median filter the results to establish a final disparity that is assigned to the grid location at the center of the filter window. If not, mark the grid point as "no disparity available".

3.4.6 Retrieve stereoscopic heights and altitude/confidence designators

3.4.6.1 Physics of the problem

In the above discussion, a cloud motion field, defined to be either a single averaged velocity vector from the Low Cloud and High Cloud bins or the two recorded velocity vectors in the two bins, was used to establish the location of the window(s) in which the search for conjugate features is performed. In this section, the cloud motion fields associated with the altitude bins serve a second purpose, namely to provide a cloud motion correction in the conversion of the disparities to heights. The methodology described in §3.4.5 results in disparities, corresponding to conjugate rays from the two pairs of cameras used in the height retrievals (ideally, An-Af and An-Aa). These conjugate rays are then intersected using a method called minimum perpendicular distance intersection. The minimum distance between the conjugate rays, and the matcher used, serve as criteria for detecting matching blunders and assigning the retrieved cloud heights with confidence levels. The difference between the forward and aft retrieved heights is also used as a blunder-detection mechanism.

The height retrieval algorithm is applied independently to the disparity fields from the reference-forward comparison and reference-aftward comparison camera pairs. If the selection between the Low Cloud and High Cloud wind assumption has not already transpired as a result of the disparity existing in a region of one of the search windows that does not overlap with the other search window, the average of the wind velocities in the two bins is used. Then, height thresholding is used to determine a near surface/above surface designation, based on the location of the feature on the SOM grid. A choice between the height associated with the reference-forward retrieval and the reference-aftward retrieval is also made.

3.4.6.2 Mathematical description of the algorithm

Due to errors associated with look directions and errors introduced during the processing of input image co-registration and image matching, two conjugate look rays rarely intersect precisely at one point. As shown in Figure 20, if $\hat{d}$ is the unit minimum distance vector between two conjugate look rays $\hat{a}$ and $\hat{b}$, then its direction is normal to both of the look rays as defined as follows:

$$\hat{d} = \begin{bmatrix} d_x \\ d_y \\ d_z \end{bmatrix} = \hat{a} \times \hat{b} = \begin{vmatrix} \hat{e}_x & \hat{e}_y & \hat{e}_z \\ a_x & a_y & a_z \\ b_x & b_y & b_z \end{vmatrix} \tag{31}$$

where $\hat{e}_x$, $\hat{e}_y$, and $\hat{e}_z$ are the unit vectors of the three axes in the Cartesian coordinate system where the look rays are intersected (i.e., the Earth Centered Reference coordinate system).

If $\vec{p}$ denotes the vector connecting the two image points P_1 and P_2 on the surface ellipsoid, then the four vectors $\vec{p}$, $\lambda_2\hat{b}$, $\lambda_1\hat{a}$, and $\lambda_0\hat{d}$ must close, where λ_0, λ_1, and λ_2 are the scale factors to the corresponding vectors. That is,

$$\vec{p} = \lambda_2\hat{b} + \lambda_0\hat{d} - \lambda_1\hat{a} \tag{32}$$

Equation (32) contains totally 3 individual equations and 3 unknowns (λ_0, λ_1, and λ_2) for each pair of conjugate rays and each intersecting point. Therefore, the intersecting object point (the midpoint of $\lambda_0\hat{d}$) can be obtained starting from the reference image location, P_2:

$$C = P_2 - \lambda_2\hat{b} - \frac{1}{2}\lambda_0\hat{d} \tag{33}$$

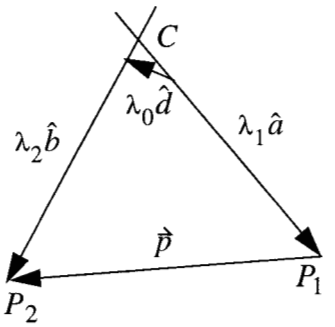

Figure 20. Object Point Intersection with minimum perpendicular distance

Now consider that the cloud moves with velocity $\vec{v_c}$ (horizontally in two dimensions) during the period Δt between the times that the conjugate cloud features were seen by the two involved cameras, where $\Delta t = t_1 - t_2$ is the observation time difference between the comparison and reference images, respectively. Note that Δt can be positive or negative. As shown in Figure 21, $\vec{a}'$ is the look vector as if there is no cloud motion, and $\vec{a}$ is the real look vector due to the existence of cloud motion $\vec{v_c}\Delta t$. If we directly intersect the two look vectors $\vec{a}$ and $\vec{b}$, the intersection will be around point D at a wrong height. To remove the cloud motion effect, we assume the cloud was actually seen by another look vector $\vec{a}'$ parallel to $\vec{a}$ at another ground position R by the forward camera at the time t_1. Then Eq. (32) becomes:

$$\vec{p} - \vec{v}\Delta t = \lambda_2\hat{b} + \lambda_0\hat{d} - \lambda_1\hat{a} \tag{34}$$

Similar to Eq. (32), there are three equations and three unknowns λ_0, λ_1, and λ_2 for each pair of conjugate rays and each intersecting point. The three-dimensional location of the cloud is obtained with Eq. (33) after the minimum distance intersection with motion is done using Eq. (34). The true

location on the SOM grid is found from the 3-dimensional Earth Centered Reference (ECR) coordinates; see *[7]*.

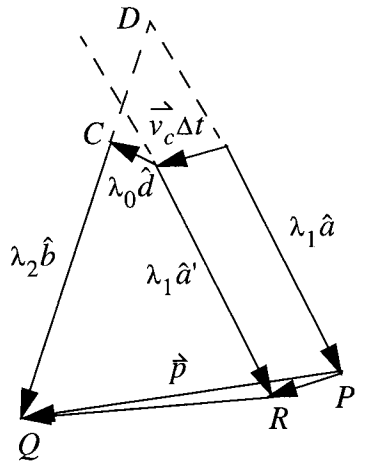

Figure 21. Minimum Distance Intersection with motion

If the disparity being converted to height (determined through application of M2, M3, or RS) is obtained within a region of the search window that is associated uniquely with either the Low Cloud bin or the High Cloud bin, the other bin is automatically eliminated and no further consideration is needed. Otherwise, if the disparity corresponds to a value that falls within the intersection of the two search windows, it is necessary to establish a method of choosing which wind velocity to use in applying the above equations. The method used is simply to average the velocities from the two bins and use this value.

For each retrieved height value, we assign a matcher score m as follows:

(1) $m = 0$ if no disparity is available or if the retrieved height is not in the range $0 \leq h \leq 20$ km.

(2) $m = 1$ if RS was used successfully and the retrieved height is in the range $0 \leq h \leq 20$ km.

(3) $m = 2$ if M3 was used successfully and the retrieved height is in the range $0 \leq h \leq 20$ km.

(4) $m = 3$ if M2 was used successfully, the final disparity was not verified to be a valid match by M3, and the retrieved height is in the range $0 \leq h \leq 20$ km.

(5) $m = 4$ if M2 was used successfully, the final disparity was also verified to be a valid match by M3, and the retrieved height is in the range $0 \leq h \leq 20$ km.

Then, for each height value, we assign a ray skewness score r as follows:

(1) If the absolute value of the ray skewness value associated with a height retrieval is > 1500 m, the result of the skewness test is designated as Blunder, and given a score of $r = 0$.

(2) If the absolute value of the ray skewness value associated with a height retrieval is < 275 m, the result of the skewness test is designated as High Confidence, and given a score of $r = 2$.

(3) Otherwise, the result of the skewness test is designated as Low Confidence, and given a score of $r = 1$.

In the above, the ray skewness is the parameter λ_0 shown in Figure 21.

The scores are combined to establish an overall score S equal to the product of m and r for the camera pair that was used for the height retrieval. We then take advantage of the fact that we have (in most cases) two independent height retrievals from the Af and Aa cameras to perform a blunder-detection test. For each domain, we compile a list of the differences in retrieved heights between the two cameras for all those $(1.1\text{-km})^2$ subregions where $S > 0$ for both forward and aft views, and then calculate the mean $\overline{\Delta h}$ and standard deviation $\sigma_{\Delta h}$ of this distribution. We then step through this domain again and for each point where there were two valid heights, determine whether the differences in the heights passes the blunder test, i.e. $\Delta h \geq \overline{\Delta h} - N_\sigma \cdot \sigma_{\Delta h}$ and $\Delta h \geq \overline{\Delta h} + N_\sigma \cdot \sigma_{\Delta h}$, where N_σ is a configurable parameter, initially set to 2.0. If the above conditions are satisfied, we set h_{ster} equal to the higher of the two heights. Else, no height is assigned to this subregion. If there is only one height with $S > 0$, we assign this height to h_{ster}.

Because the reference camera has a finite angular field-of-view (e.g., the An camera has a center-to-edge field-of-view of 15°), and because it will also have a tilt if it is not An, the retrieved height values are not necessarily located directly above the surface ellipsoid reference camera image location. As shown in Figure 22, the stereo matching on the location of E in the surface ellipsoid projected image will result in a cloud intersecting at location C. The cloud feature at location C, however, is assigned to location F on the standard Space-Oblique Mercator grid. It is at location F that the retrieved height is compared to the surface elevation to assign an altitude designation to the feature. Based on the ECR to SOM coordinate conversion described above, the feature is then assigned to the nearest corresponding point on the standard SOM grid.

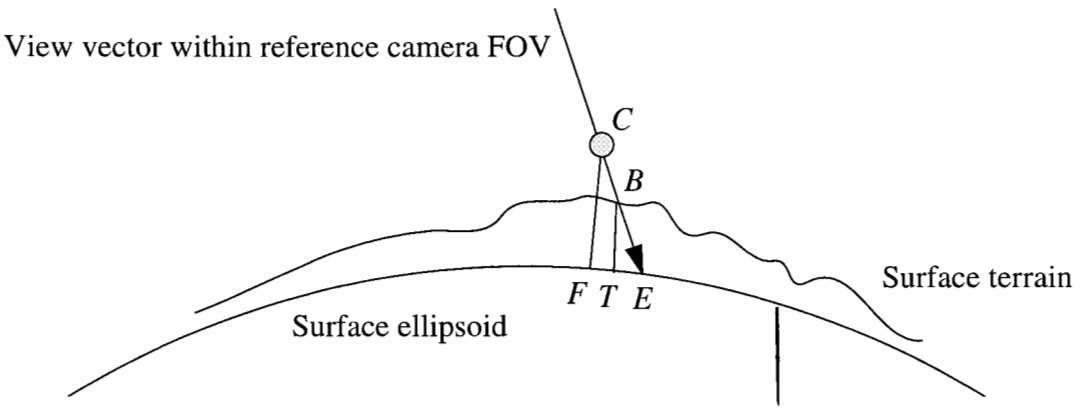

Figure 22. Projection on surface ellipsoid of reference camera view vector

We now compare the height h_{ster} to a surface threshold altitude h_s. If $h_{ster} \leq h_s$ the feature is designated NearSurf; else it is designated AboveSurf. The height threshold h_s is calculated as follows:

(1) Determine in which 1.1-km SOM subregion the height value of the feature falls, as described above; that is, establish the location of point F.

(2) Obtain from the Ancillary Geographic Product the average surface elevation at point F, $h_{surf(F)}$ and the standard deviation of the surface elevation $\sigma_{surf(F)}$.

(3) Set $h_s = h_{surf(F)} + \sigma_{surf(F)} + h_{unc}$, where h_{unc} is an additional increment to account for stereophotogrammetric uncertainties. We set $h_{unc} = 275$ m.

Additionally, if the numerical score, S, (equal to the product of m and r) is in the range 6 - 8, the height is designated High Confidence (HC). If the score is 0, it is declared No Retrieval (NR). Otherwise, it is designated Low Confidence (LC).

The result of this process is that h_{ster} has one of the following altitude/confidence designations associated with it: AboveSurfHC, AboveSurfLC, NearSurfLC, NearSurfHC, or NR.

Note that h_{ster} tells us the heights of stereoscopically matched features, but at this stage of the processing they are assigned to the geographic locations, on the SOM grid, where the reference camera view intersects the WGS84 ellipsoid (location E in Figure 22). In §3.4.8, the stereoscopic heights are assigned to the SOM grid locations which lie directly below the features (location F in Figure 22), that is, the heights are reassigned from an "ellipsoid projection" to a "feature projection".

3.4.7 Project reference camera RCCM to surface ellipsoid

3.4.7.1 Physics of the problem

We now combine the RCCM (generated during Level 1 processing) associated with the height retrieval reference camera with the stereo-matched heights in order to combine both cloud detection methods together in an objective way. Before this can be done, we must take into account the fact that the data used to derive the RCCM are projected to the surface terrain, and a ray-casting step is needed to project the RCCM onto the ellipsoid, where the reference camera image is defined. This is required because the two projections are not coincident, as shown in Figure 22. Here, a cloud detected by the RCCM projects to location B on the surface terrain, and is therefore assigned to location T on the SOM grid, and we need to project it to location E on the ellipsoid. The methodology for accomplishing this is by forward intersection of a look vector from its known location to the surface ellipsoid as described below. Additionally, it is necessary to know whether the terrain intersection point is snow or ice covered. The reason for this is that this information is used during cloud classification, as described in §3.5.5. Thus, the snow/ice indicator is carried along with the RCCM so that it is not necessary to repeat the reprojection later on in the processing.

The snow/ice indicator can be of several forms. In early versions of the data supplied by NSIDC, and in the TASC Dataset, the indicator is simply a Snow Covered/Not Snow Covered or Sea Ice Covered/Not Sea Ice Covered mask. In later versions of the NSIDC product, snow will be snow water equivalent in mm averaged over the grid cell (nominally 25 km), with values typically in the range 0 - 1000 mm, and sea ice will be coverage in percent averaged over the cell. We will interpret snow water equivalent as follows: If snow water equivalent ≥ 5 mm, then Snow Covered; else Not Snow Covered. We will interpret sea ice fraction as follows: If sea ice fraction $\geq 5\%$, then Sea Ice Covered; else Not Sea Ice Covered.

Prior to reprojection of the RCCM, it is necessary to check the glitter flag associated with the RCCM *[M-5]*. For the algorithms described in this document, we treat any glitter-contaminated RCCM as No Retrieval.

3.4.7.2 Mathematical description of the algorithm

We describe in this section the general equations to find the location E on an ellipsoid at an altitude h above the WGS84 surface ellipsoid, as seen by the same look ray passing a known topographic surface location B. For the particular problem being considered here, $h = 0$; that is, we are projecting onto the surface ellipsoid. However, we will in later sections be using similar methodology to project from the surface terrain to ellipsoids located at altitudes above the terrain, so we present the derivation in the general case. In Figure 23, B is the location on the terrain-projected surface seen by a camera with look direction described by a unit vector $\hat{l}$, E is the location of the same look ray intersecting with the ellipsoid, $\vec{p}$ and $\vec{e}$ are vectors from the center of the Earth to

the locations B and E, respectively, and finally, a and b are the semi-axes of the ellipsoid Earth.

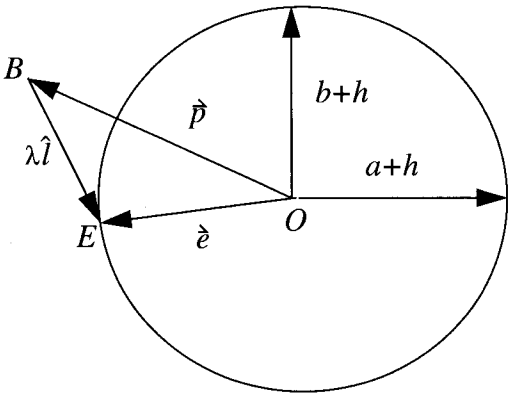

Figure 23. Intersection of a look vector with an ellipsoid

To intersect $\lambda\hat{l}$ with the ellipsoid, we first scale the ellipsoid to a unit circle. This gives us:

$$\vec{p}' = \begin{bmatrix} \frac{p_x}{a+h} \\ \frac{p_y}{a+h} \\ \frac{p_z}{b+h} \end{bmatrix} \qquad \vec{l}' = \begin{bmatrix} \frac{l_x}{a+h} \\ \frac{l_y}{a+h} \\ \frac{l_z}{b+h} \end{bmatrix} \qquad \vec{e}' = \begin{bmatrix} \frac{e_x}{a+h} \\ \frac{e_y}{a+h} \\ \frac{e_z}{b+h} \end{bmatrix} \tag{35}$$

Then we have

$$\vec{e}' = \vec{p}' \pm \lambda\vec{l}' \tag{36}$$

$$\vec{e} = \vec{p} \pm \lambda\vec{l} \tag{36a}$$

where the + sign is chosen in Eqs. (36) and (36a) if the ellipsoid lies below the terrain, and the - sign is chosen if the ellipsoid lies above the terrain. Then, since $\vec{e}'$ lies on the unit circle as shown in Figure 24, we have the condition

$$\|\vec{e}'\|^2 = \|\vec{p}' \pm \lambda\vec{l}'\|^2 = \|\vec{p}'\|^2 + \lambda^2\|\vec{l}'\|^2 \pm 2\lambda\vec{p}' \cdot \vec{l}' = 1 \tag{37}$$

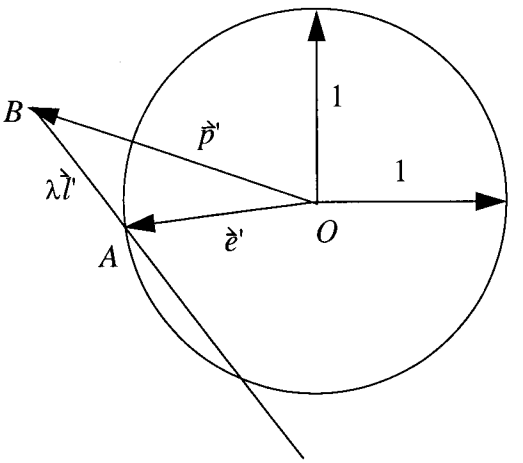

Figure 24. Two intersections of look vector with ellipsoid

Applying the quadratic formula, and choosing the appropriate root, we get

$$\lambda_{-} = \frac{-\vec{p}' \cdot \vec{l}' - \sqrt{(\vec{p}' \cdot \vec{l}')^2 + \|\vec{l}'\|^2 (1 - \|\vec{p}'\|^2)}}{\|\vec{l}'\|^2} \tag{38}$$

in the case where the ellipsoid at altitude h is below the surface terrain, and

$$\lambda_{+} = \frac{\vec{p}' \cdot \vec{l}' + \sqrt{(\vec{p}' \cdot \vec{l}')^2 + \|\vec{l}'\|^2 (1 - \|\vec{p}'\|^2)}}{\|\vec{l}'\|^2} \tag{39}$$

for the case where the ellipsoid is above the surface terrain. For the present problem, in which we are projecting the RCCM onto the WGS84 ellipsoid, we substitute the value of λ represented by Eq. (38) into Eq. (36) and choose the + sign in Eq. (36a). We then convert the Earth Centered Reference coordinates of this point to an SOM location on the ellipsoid *[7]*. This processing is done to every point in the reference camera surface-terrain-projected RCCM. Note that since the look-vectors are ellipsoid-referenced and we are projecting terrain referenced data an iterative reprojection scheme must be applied. We first use the look-vectors at the 1.1-km subregion under consideration to project down to the surface ellipsoid and then use the look-vectors at this new point to perform a new reprojection. The process stops when the maximum number of iterations is exceeded, the projected point lies beyond the bounds of the swath, or the location of the ellipsoid-projected point doesn't change between successive iterations. If the projection results in an "out-of-range" error the projection is terminated with an error message and the terrain-projected RCCM value is not assigned to any pixel on the ellipsoid. Otherwise, the last value of the ellipsoid-projected point is chosen. The forward projection could result in the following situations:

(1) A single value of the RCCM is contained within a 1.1-km subregion on the ellipsoid. The following information is carried along with this remapping for use during later processing:

(a) The value of the RCCM;

(b) The maximum terrain altitude in the "neighborhood" of T (defined below in §3.4.8.2), where T represents the SOM coordinates of point B;

(c) The SOM location and terrain altitude at T;

(d) A snow/ice indicator, established as follows: If point B is snow or ice covered; then associate the ellipsoid-projected RCCM with a Snow/Ice Covered designation. If point B is not snow or ice covered; then associate the ellipsoid-projected RCCM with a Not Snow/Ice Covered designation. The source information for the snow/ice indicator is also recorded.

(2) More than one value of the RCCM is contained within the 1.1-km subregion. In this case, the following information is carried along with this remapping for use during later processing:

(a) A single RCCM value, established as follows: Assign numerical values of 0, 1, 2, 3 and 4 respectively to the conditions No Retrieval, CloudHC, CloudLC, ClearLC, and ClearHC for all RCCM's that mapped to the subregion. The minimum numerical value is then used to select the designation for the projected RCCM such that "cloud" overrides "no-cloud";

(b) The largest of the maximum terrain altitudes in the "neighborhood" of the different T's of all points B that mapped to this subregion;

(c) SOM location of the largest of the terrain altitudes at the different T's of all points that mapped to this subregion;

(d) The largest of all the terrain altitudes T that mapped to this subregion;

(e) A snow/ice indicator, established as follows: If any of the points B are snow or ice covered; then associate the ellipsoid-projected RCCM with a Snow/Ice Covered designation. If none of the points B are snow or ice covered; then associate the ellipsoid-projected RCCM with a Not Snow/Ice Covered designation. The source information for the snow/ice indicator is also retained.

(3) No value of the RCCM is contained within the 1.1-km subregion. In this case the projected RCCM is designated No Retrieval, and no additional terrain or snow/ice designations are associated with it.

3.4.8 Assign parameters to SOM grid

3.4.8.1 Physics of the problem

Each preliminary retrieved stereoscopic height, h_{ster}, and its associated altitude/confidence designator is first compared with the ellipsoid-projected reference camera RCCM as a vehicle for overriding the retrieved heights in cases where the clear/cloud designation from the RCCM is be-

lieved to be more reliable. The output of this process is designated $\tilde{h}_{ster}$.

Next, the stereoscopic height field (the $\tilde{h}_{ster}$ values) and associated confidence levels are assigned to locations on the SOM grid. If the value of $\tilde{h}_{ster}$ exceeds the maximum cloud-top height, its confidence level is changed to "No Retrieval" when using Tables 5 and 6. These are the SOM locations of the points lying underneath the height values, that is, location F in Figure 22 (i.e., not the locations of the ray intersections with either the terrain or the ellipsoid). We also assign to these SOM grid locations the values of the ellipsoid-projected reference camera RCCM and the associated snow/ice indicator, as well as the retrieved (preliminary) stereoscopic height, h_{ster}.

Finally, for those SOM grid locations where (a) a default cloud height was assigned, (b) the MODIS cloud mask indicates the presence of cloud, and (c) MODIS cloud height data are available, the default cloud height is further overridden by the MODIS cloud height at the SOM grid point. Note that when this is done, the location on the SOM grid where the RCCM is assigned still corresponds to the placement determined by using the MISR default cloud height. A readjustment of its placement would make the algorithm more complex. The reason that we do not immediately assign the MODIS height, rather than a default, is that the MODIS look vectors are not aligned with the MISR RCCM look vectors. We opt not to incur additional complexity. Consequently, the SOM-projected RCCM may be mislocated in these instances. However, since the only purpose of the SOM-projected RCCM is for regional cloud classification over $(17.6\text{-km})^2$ regions (see §3.5.5), the effect of this mislocation on the regional classifiers should be small.

3.4.8.2 Mathematical description of the algorithm

The decision matrix for determining whether or not to override the stereoscopic heights is shown in Table 5. If the value of $\tilde{h}_{ster}$ exceeds the maximum cloud-top height, its confidence level is changed to "No Retrieval" when using Tables 5 and 6, and the reason-for-override flag is set to "Cloud Too High".

Table 5: Decision matrix for establishing final stereoscopic heights $\tilde{h}_{ster}$

		Ellipsoid-projected reference camera RCCM				
		No Retrieval	CloudHC	CloudLC	ClearLC	ClearHC
h_{ster} designator	No Retrieval	No Retrieval	Default Cloud	Default Cloud	Surface	Surface
	AboveSurfHC	Keep Stereo	Keep Stereo	Keep Stereo	Keep Stereo	Surface
	AboveSurfLC	Keep Stereo	Keep Stereo	Keep Stereo	Surface	Surface
	NearSurfLC	Keep Stereo	Keep Stereo	Keep Stereo	Surface	Surface
	NearSurfHC	Keep Stereo	Keep Stereo	Keep Stereo	Surface	Surface

Three possibilities can occur, as indicated in Table 5:

(1) The stereoscopically retrieved height is kept, and we set $\tilde{h}_{ster} = h_{ster}$. This situation is designated "Keep" in Table 5. Additionally, the stereo-override flag is set to "Keep Stereo".

(2) The stereo height h_{ster} is overridden and we set $\tilde{h}_{ster} = h_{surf(T)}$, the surface elevation at the SOM grid location where the reference camera view intersects the surface (i.e., point T in Figure 22, which is the SOM location of terrain point B). This is done in situations where the Clear designation of the RCCM is believed to be the more reliable indicator, and is shown as "Surface" in Table 5. If the override flag has not already been set to "Cloud Too High", it is set to "RCCM Override".

(3) When the RCCM indicates cloud and the stereo retrieval has failed, we set $\tilde{h}_{ster} = h_{cld}$, a default cloud altitude. This default is defined by $h_{cld} = max\{h_{surf(T-neighborhood)}\} + 3$ km. The first term on the right-hand-side of this expression denotes the maximum surface elevation in the "neighborhood" of point T. The definition of this term, and the rationale for choosing this approach (rather than, for example, simply using $h_{surf(T)}$), is presented below. This situation is designated "Default Cloud" in Table 5. If the override flag has not already been set to "Cloud Too High", it is set to "RCCM Override".

The assignment of SOM grid locations is accomplished by using the height retrieval reference camera look vector and the following logic:

(1) If $\tilde{h}_{ster}$ was established according to a situation designated "Default Cloud", we imagine an ellipsoid at a height $\tilde{h}_{ster}$. Using the methodology described in §3.4.7.2, projecting to an ellipsoid at an altitude $h = \tilde{h}_{ster}$ above the WGS84 ellipsoid, and choosing the case in which the ellipsoid is above the terrain, we calculate the ECR coordinates of the cloud. These are then converted to SOM grid coordinates.

(2) If $\tilde{h}_{ster}$ was established according to a situation designated "Keep Stereo", we assign the point to the SOM grid location *F*.

(3) If $\tilde{h}_{ster}$ was established according to a situation designated "Surface", we assign the point to the SOM grid location *T*.

In the nominal case in which An is the reference camera, and given the center-to-edge FOV of the An camera of 15°, choice (1) will result in at most about a 0.8 km horizontal difference between the feature location and SOM grid location *T*, and thus may put the cloud feature one 1.1-km SOM grid element away from location *T*. In the extreme case in which a B camera is the reference, the feature could be reassigned to a location that is about three 1.1-km SOM grid elements away from location *T*. We need to insure that when this occurs, the altitude is higher than the surface elevation at the new SOM location. In the situations designated "Keep Stereo" in Table 5, this is guaranteed. However, in the "Default Cloud" situation where we set $\tilde{h}_{ster} = h_{cld}$, if we simply added the default of 3 km to the surface elevation at point *T* to establish h_{cld}, we could not guarantee that this would occur, because the surface terrain is continuously varying. Thus, we have adopted the following approach. We construct a 7 x 7 array of 1.1-km SOM grid points, in which grid location *T* is at the center. This defines the "neighborhood" of location *T*. We then define $max\{h_{surf(T-neighborhood)}\}$ to be the maximum surface elevation in this 7 x 7 array of points. Using this approach and defining h_{cld} as given in §3.4.8.2 guarantee that the cloud will be at an altitude higher than the surface location of the reassigned SOM grid point.

To maintain consistency with the final stereoscopic heights, the altitude and confidence designations associated with each are replaced with the designations shown in Table 6. Thus, the initial designations No Retrieval, AboveSurfHC, AboveSurfLC, NearSurfLC, and NearSurfHC are replaced, as indicated, with No Retrieval, CloudHC, CloudLC, NearSurface, and Clear.

Table 6: Classifications associated with the final stereoscopic heights $\tilde{h}_{ster}$

		Ellipsoid-projected reference camera RCCM				
		No Retrieval	CloudHC	CloudLC	ClearLC	ClearHC
h_{ster} designator	No Retrieval	No Retrieval	CloudLC	CloudLC	Clear	Clear
	AboveSurfHC	CloudHC	CloudHC	CloudHC	CloudHC	Clear
	AboveSurfLC	CloudLC	CloudLC	CloudLC	Clear	Clear
	NearSurfLC	NearSurface	NearSurface	NearSurface	Clear	Clear
	NearSurfHC	NearSurface	NearSurface	NearSurface	Clear	Clear

When following the above logic and assigning the stereoscopic height field to the SOM grid, the following situations can occur:

(1) A single stereoscopic height is contained within a 1.1-km subregion on the SOM ellipsoid. In this case there are three possibilities:

(a) This height resulted from the decision "Keep Stereo" in Table 5. The final SOM-projected height field is assigned a designation based on the altitude and confidence designation associated with it. Additionally, a flag value specifying that this height was kept from the original stereoscopic retrieval is established.

(b) This height resulted from the decision "Keep Stereo" in Table 5, but its value is lower than the surface height of the SOM grid cell within which it falls. In this case we override the value with the surface height, and assign the altitude/confidence designation Clear. A flag value specifying that this height was overridden with the surface height for this reason is established. The stereo-override flag is set to "Cloud Too Low".

(b) This height resulted from an override decision in Table 5, that is, the preliminary stereoscopic value was changed to "Surface" or "Default Cloud". The altitude/confidence designation associated with this value is also assigned to the 1.1-km subregion, along with a flag specifying that this was a "Surface" override or "Default Cloud" height from Table 5, as appropriate. If a MODIS cloud height is available, we replace the modified stereoscopic height by the MODIS height, but only if the MODIS altitude is higher than the modified stereoscopic height. This condition also requires the MODIS cloud mask to indicate the presence of cloud. When this occurs, the altitude/confidence designation is specified to be CloudLC, and an additional flag indicating that the source was MODIS is set.

(2) More than one stereoscopic height is contained within the 1.1-km subregion. In this case there are five possibilities:

(a) All of these heights resulted from the decision "Keep Stereo" in Table 5.We choose the maximum height value, and the altitude/confidence designation associated with it, and assign this to the 1.1-km subregion, along with a flag specifying that this was a stereoscopically derived height. In the case of equal heights but unequal altitude/confidence designators, we choose the value of the altitude/confidence flag which is closest to CloudHC.

(b) All of these heights resulted from the decision "Keep Stereo" in Table 5 and the maximum height value is chosen, but this value is lower than the surface height of the SOM grid cell within which it falls. In this case we override the value with the surface height, and assign the altitude/confidence designation Clear. A flag value specifying that this height was overridden with the surface height for this reason is established and the stereo-override flag is set to "Cloud Too Low".

(c) The available heights consist of a mixture of stereoscopically retrieved heights and modified heights ("Surface" or "Default Cloud"). In this case the modified heights are excluded from the computation of the maximum value, and the designations associated with this height are specified as in (a).

(d) The logic of (c) is followed, and the resulting height is lower than the surface height of the SOM grid cell within which it falls. In this case we override the value with the surface height, and the designations associated with this height are specified as in (b).

(e) All of these heights resulted from an override decision in Table 5, that is, the preliminary stereoscopic values were changed to "Surface" or "Default Cloud". Again, we choose the maximum height value, and the altitude/confidence designation associated with it, and assign this to the 1.1-km subregion, along with a flag specifying that this was a "Surface" override or "Default Cloud" height, as appropriate. As above, if we have equal heights but unequal altitude/confidence designators, we choose the value of the altitude/confidence flag which is closest to CloudHC. If a MODIS cloud height is available, we replace this modified stereoscopic height by the MODIS height, but only if the MODIS altitude is higher than the modified stereoscopic height. This condition also requires the MODIS cloud mask to indicate the presence of cloud. When this occurs, the altitude/confidence designation is specified to be CloudLC, and an additional flag indicating that the source was MODIS is set.

(3) No stereoscopic height is contained within the 1.1-km subregion. In this case there are two possibilities:

(a) The MODIS cloud height, if available, is assigned. This requires the MODIS cloud mask to indicate the presence of cloud. When this occurs, the altitude/confidence designation is specified to be CloudLC, and an additional flag indicating that the source was MODIS is set.

(b) MODIS cloud height is not available. In this case the SDCM and the height source flag indicate No Retrieval.

There are two separate flags that have to be kept -- one specifying the source of the final stereoscopic height and the other the reason for the override of the preliminary stereoscopic height (if it occurred). Both of these flags have to be feature-projected to the SOM grid along with the value of the final stereoscopic height ($\tilde{h}_{ster}$)and the SDCM.

Note that the value of h_{ster} is not assigned to the same feature-projected SOM location as $\tilde{h}_{ster}$, but is rather maintained at its original feature-projected location (before possibly being overwritten by a surface height).

The SDCM is established for each 1.1-km subregion on the SOM grid. The SDCM is equal to the altitude/confidence designators that are associated with the feature-projected $\tilde{h}_{ster}$ field, and can take on one of five values: CloudHC, CloudLC, NearSurface, Clear, and No Retrieval.

When reassigning the ellipsoid-projected RCCM and associated snow/ice cover designation to the feature-projected SOM coordinates, the following situations can occur:

(1) A single value of the RCCM is contained within a 1.1-km subregion on the SOM grid. The feature-projected RCCM is assigned this value. The snow/ice designation associated with the RCCM is also assigned to this SOM location to establish the feature-projected snow/ice designation. The snow/ice source flag is also kept.

(2) More than one non-No Retrieval value of the RCCM is contained within the 1.1-km subregion. In this case we do the following:

(a) Assign numerical values of 1, 2, 3 and 4 respectively to the conditions CloudHC, CloudLC, ClearLC, and ClearHC. The minimum numerical value is then used to select the projected RCCM designation;

(b) Determine whether any of the RCCM's which map to the subregion has a Snow/Ice Covered designation associated with it. If this is the case, we establish this subregion as having a feature-projected Snow/Ice Covered designation. Otherwise, we establish the subregion as having a feature-projected Not Snow/Ice Covered designation. The snow/ice source information is also retained. In the case of equal snow/ice values, we choose the source that is deemed to be the most reliable (MODIS, then NSIDC, then DAO, and finally TASC) and set the flag accordingly.

(3) No value of the RCCM is contained within the 1.1-km subregion. In this case the projected RCCM is designated No Retrieval, and no associated snow/ice information is carried along.

Note that the feature-projected snow/ice designation indicates the presence or absence of snow/ice in the line-of-sight of the reference camera that intersects the feature at height $\tilde{h}_{ster}$ above the ellipsoid at this location.

3.4.9 Determine RLRA and RLRA standard deviation

3.4.9.1 Physics of the problem

We now use the final, feature-projected stereoscopic heights $\tilde{h}_{ster}$ to determine values of the RLRA and the standard deviation of RLRA.

Within any given region we expect there to be differences in the cloud top height. These differences may on occasion be negligible for very stratiform clouds, about 100 m for stratocumulus, and on the order of kilometers for broken cumuliform cloud fields or multilayered systems with maximum contrast at different levels within the region. For the purposes of albedo calculations *[M-9]*, the way we reproject the data and add up radiances corresponding to a specific set of RLRA's requires us to take in to careful account the effect of one region obscuring the line of sight to another region with a lower RLRA (see §3.4.10). For Level 2 standard products, we do not account for the recapture of lines of sight underneath clouds (all potentially obscured radiances are binned into a category labelled "side-leaving radiances" and not processed in terms of high resolution albedos), so for consistency all we need to know is whether there is any possibility that one region obscures the line of sight to another region. This is determined by the maximum altitude of a feature within the 2.2-km subregion within which the RLRA is defined, and thus is a better choice than setting the RLRA to some average altitude. This philosophy is reflected in the way the RLRA is generated from the stereoscopic heights.

3.4.9.2 Mathematical description of the algorithm

3.4.9.2.1 Determine RLRA

Each 2.2-km subregion on the SOM grid will typically have four values of $\tilde{h}_{ster}$ projected into it; however, the projection process may result in certain 1.1-km subregions having no features assigned. Consequently, the composite 2.2-km subregion may have less than four values. Whatever the number of height values that are present, the maximum value within each 2.2-km subregion defines the RLRA. The following logic is used:

(1) If the set of $\tilde{h}_{ster}$ values consists entirely of heights based on original (unmodified) stereoscopic retrievals, or alternatively, if the set of $\tilde{h}_{ster}$ values consists entirely of heights based on overrides with the default cloud altitude, MODIS cloud altitude, or surface elevation (either due to the reasoning of Table 5 or because the preliminary stereoscopic height was lower than the surface elevation), then

$$RLRA = max\{\tilde{h}_{ster}\} \tag{40}$$

(2) If the set of $\tilde{h}_{ster}$ values consists of a mixture of stereoscopically-retrieved heights and heights based on overrides with the default cloud altitude, MODIS cloud altitude, or surface elevation (either due to the reasoning of Table 5 or because the preliminary stereoscopic height was lower than the surface elevation), then the override heights are excluded from the right-hand-side of Eq. (40).

(3) If there are no $\tilde{h}_{ster}$ values available from any source, it is necessary to establish a value of RLRA so that the RLRA surface does not have any gaps. In this event we establish a set of four altitudes h_{surf}, one for each of the four 1.1-km subregions making up the 2.2-km subregion defining the RLRA. Then, the "fill" value of RLRA is

$$RLRA = max\{h_{surf}\} + 3 \text{ km} \tag{40a}$$

3.4.9.2.2 Calculate standard deviation in RLRA

The results are combined to calculate a standard deviation associated with the RLRA to give a measure of the variation among the $\tilde{h}_{ster}$ values found within the 2.2 km subregion. Thus, we let

$$\sigma_{RLRA} = \sqrt{\langle \tilde{h}_{ster}^2 \rangle - \langle \tilde{h}_{ster} \rangle^2} \tag{41}$$

where all available values of $\tilde{h}_{ster}$, except for those assigned a default cloud height, are included in the calculation. If only a single value of $\tilde{h}_{ster}$ is available, σ_{RLRA} is set to a flag indicating this condition. If no value of $\tilde{h}_{ster}$ was available, such that the RLRA was established using Eq. (40a), σ_{RLRA} is set to a flag indicating this condition. Finally, if RLRA is set to No Retrieval, σ_{RLRA} is set to a flag indicating this condition.

3.4.10 Determine RLRA projections

3.4.10.1 Physics of the problem

A $(2.2 \text{ km})^2$ RLRA sample is regarded as a square prism with its top at the RLRA and four side walls standing from the surface ellipsoid. Camera look directions are provided at the surface ellipsoid, by which we can link the surface ellipsoid to the RLRA. Because the RLRA is spatially variable, we use an approach in which we project each 2.2-km RLRA sample down to the analytical surface ellipsoid. Once the projection between the RLRA and surface ellipsoid is found, we determine whether a camera look vector intersects the top of an RLRA prism, or an along-track or cross-track side of the prism, and whether a particular camera's view is obscured by another RLRA prism.

The projection of one RLRA prism onto the surface ellipsoid is shown in Figure 25. An

RLRA prism is at 2.2 km resolution in SOM space. We assume the surface ellipsoid is represented at 275-m resolution in the SOM space. The mapping between the center location O of the RLRA prism top to the surface ellipsoid can be found by an algorithm which intersects along the look vector from point O to the surface ellipsoid at O_t.

The first step of the projection algorithm is to determine the zenith angle and azimuth angle of the projection to the ellipsoid, relative to the SOM grid. These angles are then used to project the "target" prism (the prism for which reflecting level parameters are being determined) as well as surrounding prisms which potentially obscure the target prism. These projections are performed as described below, and an ellipsoid projection map is used to keep track of where the top and sides of the target prism map to the ellipsoid, and which of the pixels within the projected area are obscured by neighboring prisms. Projection maps are created separately for each target prism and its neighbors. Once the map is complete, the results are used to generate the reflecting level parameters, and then the map is "erased" and the next target prism is considered. This process proceeds separately for each MISR camera. The red band view vector is used to establish the mappings and the results are then applied to the other three bands. For this purpose, the small difference in view zenith angle between the red and non-red bands is ignored.

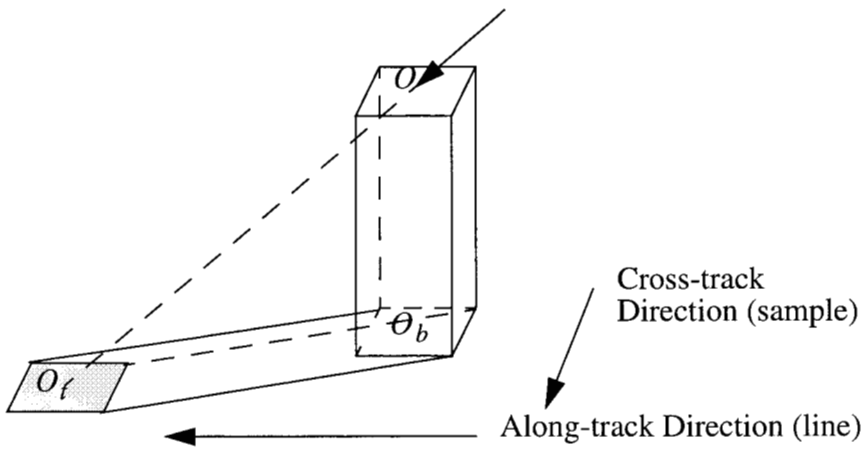

Figure 25. The projection of a RLRA prism on the surface ellipsoid

3.4.10.2 Mathematical description of the algorithm

The look vector which must be used to project the target prism as well as neighboring, potentially obscuring prisms, is the vector associated with the point O_t on the ellipsoid. Since we do not a priori know the location of O_t, the look vector at location O_b is used as a first guess. This vector is then used to project point O to the ellipsoid, providing an estimate of the location O_t. The look vector at this location is then used to reproject point O to obtain a new estimate of the location O_t. This iterative process is stopped when O_t does not change by more than 1 pixel in both the along-track and cross-track directions.

Equation deleted (42)

Once the appropriate look vector is determined, the zenith angle θ and azimuth angle ϕ of the look vector relative to the SOM coordinate system, at the surface ellipsoid, are established. Using these angles, a prism projects to the ellipsoid as depicted in Figure 26. Because of the oblique view angles, and the maximum height of an RLRA of 20 km above the surface ellipsoid, it is necessary to take into account the curvature of the surface ellipsoid in calculating the along-track and cross-track projection lengths. For this purpose, it sufficiently accurate to represent the ellipsoid locally as a sphere with a radius equal to the mean radius of the Earth, r_e. The along-track and cross-track projection lengths, respectively, are given by:

$$\Delta a = \left[\theta - \arcsin\left(\frac{r_e}{r_e + h} \cdot \sin\theta\right)\right] \cdot r_e \cdot \cos\phi$$

$$\Delta c = \left[\theta - \arcsin\left(\frac{r_e}{r_e + h} \cdot \sin\theta\right)\right] \cdot r_e \cdot \sin\phi \qquad (42a)$$

where h is the RLRA of the prism. These are converted to numbers of pixels by dividing these distances by 275 m and rounding off to the nearest pixel.

The projection map is then filled in by casting lines of projection from the edges of the prism top projection that are furthest from the base, back toward the base. Those pixels which fall within the top projection, an area measuring 8 pixels x 8 pixels (2.2 km x 2.2 km) are marked "top", and the remainder are marked "side". The number of top pixels is always 64 and the number of side pixels is established during the casting. We also initialize the number of unobscured pixels on the top, and the number of unobscured pixels on the side, to these values. These latter parameters will be updated during the process of projecting the neighboring prisms, and reduced in accordance with the obscurations that are found. Note that at the edges of the projection of the prism sides, a pixel on the projection map may only be partially included within the shaded area of Figure 26 that marks the side projection. For simplicity, we consider any pixel that is only partially included within this area to be approximated as 50% included, and contributes to the count of number of side pixels with a weight of one-half relative to the pixels that are entirely included within the projected area.

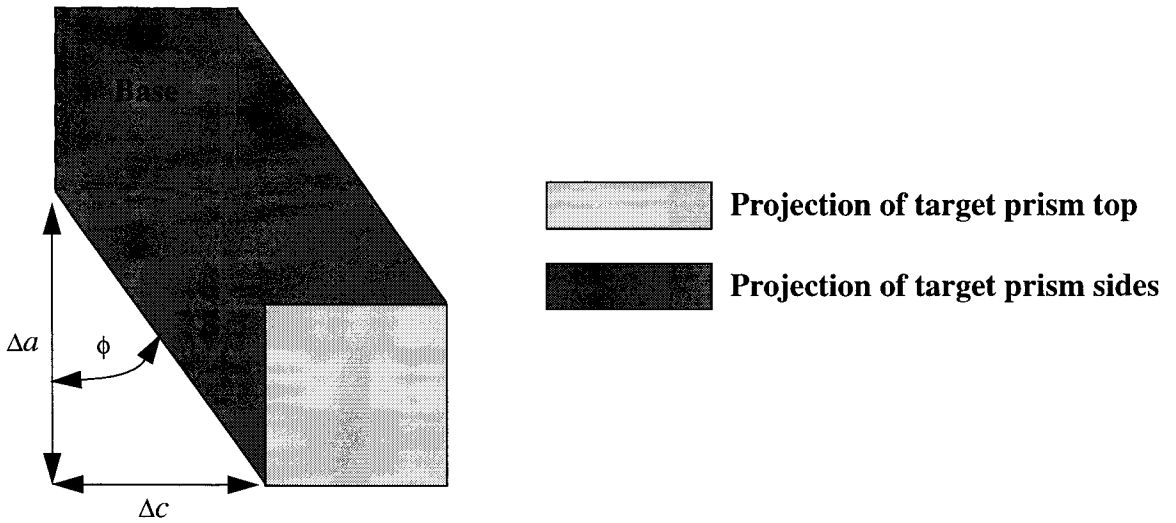

Figure 26. RLRA prism projection

A similar procedure is now applied to the prisms which are the neighbors of the target prism, and which are situated within the block so as to potentially obscure either the top or side of the target prism. This is depicted in Figure 27. Obscuration means that radiation leaving the prism top or side, and moving in the direction of the look vector, but toward the camera, is impeded by another RLRA prism before reaching the camera; in other words, the radiation detected by the camera at that look direction came from a different prism. The projections of these prisms to the projection map is treated in a similar fashion, and use the same values of zenith angle θ and azimuth angle ϕ as for the target prism. Pixels within the target prism's projected area are now marked as obscured if the neighboring prism projection overlaps them.

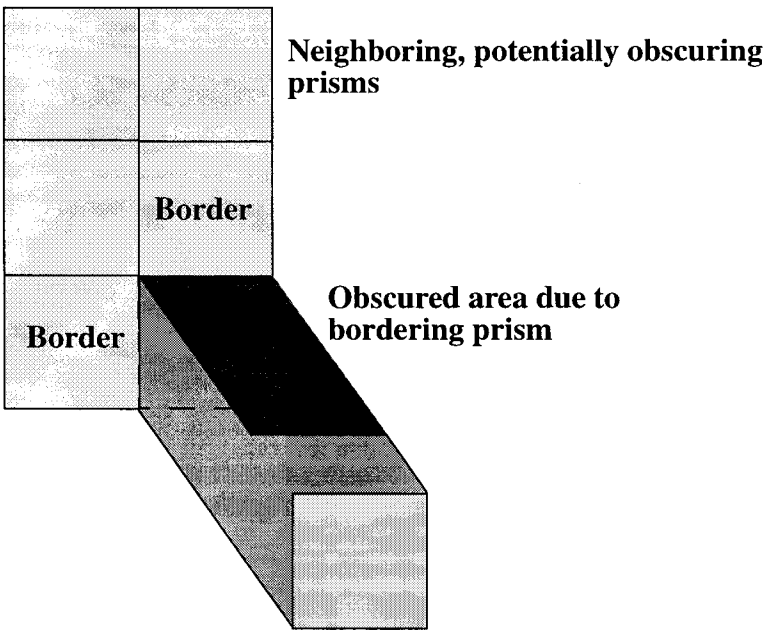

Figure 27. RLRA projection of neighboring prisms

For neighboring prisms, we consider projected pixels that fully overlap projected pixels of the target prism to be completely obscuring, while for simplicity, projected pixels that only partially overlap the projected pixels of the target prism are considered to be obscuring with a weight of 50%. The end result of this process, for a given target prism and camera angle, is a 275-m map which indicates whether a pixel contains the projection of the top or side of the target prism, and whether the pixel is unobscured, partially obscured, or completely obscured.

Once the projection of the neighboring prisms is complete for a given target prism, we then use the projection map and the BRF data at the ellipsoid to calculate the reflecting level parameters. During the calculation of the reflecting level parameters, the projection map for the target prism is simultaneously "erased", preparing the projection map for use with the next target prism.

3.4.11 Calculate reflecting level parameters

3.4.11.1 Physics of the problem

Reflecting level parameters consist of average BRF's projected to the RLRA prisms and number of unobscured pixels on the tops and sides. These are used for calculations of local albedos. Additionally, texture indices are calculated at the tops of RLRA prisms, at each of the nine MISR angles. New methods of classifying cloud scenes based on the variation of contrast (texture) as a function of view zenith angle will be investigated with MISR data. Texture is characterized by a

set of statistical measures based on band 3 (red) BRF spatial distributions. One simple measure is the ratio of the standard deviation to the mean of the 275-m red channel BRF's assigned to the top surface of a 2.2-km column at the reflecting level. This provides a measure of the scene homogeneity. Additional measures are the standard deviations of the absolute values of adjacent pixel differences in BRF, calculated separately for the cross-track and along-track directions. These provide indications of scene brokenness, and are equivalent to the Gray Level Difference Vector (GLDV) Standard Deviation *[1]*, which is reported to be one of the best discriminators between cloud classes. In summary, the reflecting level parameters on the top of an RLRA prism are the following:

(a) The average BRF values (for all four MISR spectral bands, calculated from all unobscured and partially obscured pixels which have RDQI $\leq RDQI_4$, where we set $RDQI_4 = 0$);

(b) The number of unobscured pixels (calculated for the red band only, regardless of RDQI. The reason that RDQI is immaterial here is that this is a geometric, not a radiometric indicator);

(c) Three texture indices (calculated for the red band only, using all unobscured pixels which have RDQI $\leq RDQI_4$).

Texture indices are not calculated on the prism sides. The reflecting level parameters on the side of an RLRA prism are the following:

(a) The average BRF values (for all four MISR spectral bands, calculated from all unobscured and partially obscured pixels which have RDQI $\leq RDQI_4$);

(b) The number of unobscured pixels (calculated for the red band only, regardless of RDQI. The reason that RDQI is immaterial here is that this is a geometric, not a radiometric indicator).

3.4.11.2 Mathematical description of the algorithm

Once the projection and obscuration map for a given target prism has been calculated, it is used along with the ellipsoid-projected BRF's and RDQI's to generate the top and side reflecting level parameters. The calculation of reflecting-level parameters can be accomplished simultaneously with "erasing" the projection map in preparation for usage with the next target prism. This "erasure" procedure involves marching one last time down the projection lines 2 - 16 shown in Figure 26, and resetting the value in each pixel to the value for a pristine map. This process provides the following inputs.

(a) An indication of whether the pixel is mapped to the top or side of the target prism.

(b) The value of t, the count of how many of the 64 top pixels are unobscured. This is simply rounded off to the nearest integer and reported.

(c) The value of s, the count of how many of the side pixels are unobscured. During algorithm testing, it was found that the accuracy of this parameter could be improved using the following approach. First, it is assumed that the parameter

$$f_{side} = \frac{\text{number of unobscured side pixels}}{\text{total number of side pixels}} \tag{42aa}$$

calculated from the projection algorithm is determined with higher accuracy than the numerator of Eq. (42aa) alone. We then let

$$s = f_{side} \cdot s_{tot} \tag{42ab}$$

where s_{tot} is an analytic estimate of the total number of side pixels, given by

$$s_{tot} = nint\left\{8 \cdot \frac{[|\Delta a| + |\Delta c|]}{0.275}\right\} \tag{42ac}$$

(d) w_{ij}, the obscuration weight associated with the i,j^{th} pixel. We assume a value of 0 for obscured pixels, 1 for partially obscured pixels, and 2 for unobscured pixels. These weights are used in the calculation of average BRF's.

(e) w'_{ij}, a modified obscuration weight derived from w_{ij}, defined to be 0 for obscured or partially obscured pixels, and 1 for unobscured pixels. The modified weights are used in the calculation of texture indices, which make use of only completely unobscured pixels.

(f) $B_{ij}(\lambda)$, the ellipsoid-projected BRF associated with the i,j^{th} pixel. For non-nadir cameras and non-red band channels, the surface ellipsoid data are provided in 1 x 4 or 4 x 4 resolution. When computing the reflecting level parameters for these channels, we first reduce the resolution to 1.1 km if necessary. This is required in order to make the appropriate assignments of the BRF values to the RLRA tops and sides.

(g) v_{ij}, a data validity weight assigned a value of 1 if the ellipsoid RDQI of the i,j^{th} pixel $\leq$ $RDQI_4$, and 0 otherwise.

(h) a_{ij}, an along-track weight assigned a value of 1 if the ellipsoid RDQI of the i,j^{th} pixel $\leq$ $RDQI_4$ and if the ellipsoid RDQI of the adjacent pixel to i,j in the along-track direction $\leq$ $RDQI_4$, and if both the i,j^{th} pixel and the adjacent pixel are labeled "top". Otherwise the weight is assigned a value of 0. The adjacent pixel is the one closer to the corner of the target prism at which the two border prisms meet; this one is used since its obscuration weights have not yet been "erased" (see Figure 26).

(i) $B_{adjacent,\,along}(red)$, the red band BRF in the adjacent pixel defined in (h).

(j) c_{ij}, a cross-track weight assigned a value of 1 if the ellipsoid RDQI of the i,j^{th} pixel $\leq \text{RDQI}_4$ and if the ellipsoid RDQI of the adjacent pixel to i,j in the cross-track direction $\leq \text{RDQI}_4$, and if both the i,j^{th} pixel and the adjacent pixel are labeled "top". Otherwise the weight is assigned a value of 0. The adjacent pixel is the one further from the corner of the target prism at which the two border prisms meet; this one is used since its obscuration weights have not yet been "erased" (see Figure 26).

(k) $B_{adjacent,\,cross}(red)$, the red band BRF in the adjacent pixel defined in (j).

At the same time as the "erasure" of the projection map occurs, the sums required to perform the following calculations are accumulated, and the reflecting-level parameters are then determined.

(1) The average top-projected BRF in each band, $\bar{B}_{top}(\lambda)$, is calculated from:

$$\bar{B}_{top}(\lambda) = \frac{\sum_i \sum_j v_{ij} w_{ij} B_{ij}(\lambda)}{\sum_i \sum_j v_{ij} w_{ij}} \tag{42b}$$

where the required sums are accumulated only over those pixels that are marked "top".

(2) The first top texture index is determined as follows. Calculate the mean of the squared BRF's in the red band from:

$$\langle r^2 \rangle = \frac{\sum_i \sum_j v_{ij} w_{ij} (B_{ij,\,red})^2}{\sum_i \sum_j v_{ij} w_{ij}} \tag{42c}$$

where the required sums are accumulated only over those pixels that are marked "top". Let $\langle r \rangle$ be the mean value of the red band BRF calculated using Eq. (42b). Then compute the standard deviation from:

$$\sigma = \sqrt{\langle r^2 \rangle - \langle r \rangle^2} \tag{43}$$

The texture index is then given by $\zeta = \sigma / \langle r \rangle$.

(3) The second top texture index, the along-track GLDV standard deviation, is calculated as follows. Calculate the mean red-band BRF along-track absolute difference from:

$$\langle \Delta r_{along} \rangle = \frac{\sum_i \sum_j a_{ij} w'_{ij} \left| B_{ij}(red) - B_{adjacent,\,along}(red) \right|}{\sum_i \sum_j a_{ij} w'_{ij}} \tag{43a}$$

and the mean red-band BRF along-track absolute difference squared from:

$$\langle \Delta r^2_{along} \rangle = \frac{\sum_i \sum_j a_{ij} w'_{ij} \left| B_{ij}(red) - B_{adjacent,\,along}(red) \right|^2}{\sum_i \sum_j a_{ij} w'_{ij}} \tag{43b}$$

where the required sums are accumulated only over those pixels that are marked "top". The required texture index σ_{along} is determined from

$$\sigma_{along} = \sqrt{\langle \Delta r^2_{along} \rangle - \langle \Delta r_{along} \rangle^2} \tag{43c}$$

(4) The third top texture index, the cross-track GLDV standard deviation, is calculated as follows. Calculate the mean red-band BRF cross-track absolute difference from:

$$\langle \Delta r_{cross} \rangle = \frac{\sum_i \sum_j c_{ij} w'_{ij} \left| B_{ij}(red) - B_{adjacent,\,cross}(red) \right|}{\sum_i \sum_j c_{ij} w'_{ij}} \tag{43d}$$

and the mean red-band BRF along-track absolute difference squared from:

$$\langle \Delta r^2_{cross} \rangle = \frac{\sum_i \sum_j c_{ij} w'_{ij} \left| B_{ij}(red) - B_{adjacent,\,cross}(red) \right|^2}{\sum_i \sum_j c_{ij} w'_{ij}} \tag{43e}$$

where the required sums are accumulated only over those pixels that are marked "top". The required texture index σ_{cross} is determined from

$$\sigma_{cross} = \sqrt{\langle \Delta r^2_{cross} \rangle - \langle \Delta r_{cross} \rangle^2} \tag{43f}$$

(5) The average side-projected BRF in each band, $\bar{B}_{side}(\lambda)$, is calculated from:

$$\bar{B}_{side}(\lambda) = \frac{\sum_i \sum_j v_{ij} w_{ij} B_{ij}(\lambda)}{\sum_i \sum_j v_{ij} w_{ij}} \tag{43g}$$

where the required sums are accumulated only over those pixels that are marked "side".

3.5 THEORETICAL DESCRIPTION: STAGE 3 PROCESSING

In Stage 3, regional (17.6-km) cloud classification occurs. The first step is application of a cirrus cloud detection method to generate a cloud mask known as the Angular Signature Cloud Mask (ASCM). The results of this process are then combined with the SDCM and RCCM to generate the regional classifiers. Figure 6 depicts a flow-diagram of the algorithms required for Stage 3.

3.5.1 Perform BDAS test

3.5.1.1 Physics of the problem

High-altitude cirrus clouds will often lie below the thresholds set in the RCCM production performed in Level 1B2, such that cirrus-containing areas may be classified as clear. To detect cirrus, we will make use of MISR's unique multi-angle viewing capabilities through a technique known as Band-Differenced Angular Signature, or BDAS *[2]*, *[3]*. The result of this process, after projection to the standard SOM grid, is the generation of an Angular Signature Cloud Mask, or ASCM.

The BDAS technique uses the difference in BRF between either bands 1 and 4 or bands 1 and 3, depending on surface type, to detect cirrus cloud presence. The basis of the technique is that the BDAS is sensitive to the relative contribution of Rayleigh scattering to the total reflectance. Bands 1 and 4 differ by almost a factor of 15 in their sensitivity to Rayleigh scattering, while bands 1 and 3 differ by more than a factor 5. Since the contribution of Rayleigh scattering decreases in the presence of high clouds, the resulting changes in BDAS enable detection of such clouds. Simulations using LOWTRAN 7 indicate that the technique is expected to work best over ocean and, interestingly, over snow surfaces (where most other satellite cloud detection techniques experience difficulties). Over most land surfaces, the BDAS technique generally offers poor cloud detection capability due to the large difference in land reflectance between bands 1 and 4; over these surfaces, therefore, the band 1 - 3 difference is used. This has the advantage of reducing the difference in surface spectral reflectance between the two bands; however, it has the disadvantage of reducing the difference in the Rayleigh scattering contribution between the two bands, thus degrading its sensitivity at detecting very thin cirrus.

Di Girolamo and Davies *[3]* have shown that a single, global BDAS-threshold can be used for detecting cirrus; however, additional scene information would ultimately lead to better BDAS cloud detectability, especially over land. For this reason, the Angular Signature Threshold Dataset will be updated after launch.

3.5.1.2 Mathematical description of the algorithm

All cloud detection algorithms require a threshold that distinguishes clear and cloudy populations. The BDAS technique is no different. However, the absolute value of the band-differenced reflectance factor (which depends on cosine of the view zenith angle, μ, the cosine of the solar zenith angle μ_0, and the view-solar azimuth difference ϕ - ϕ_0), defined as

$$\Delta r(-\mu, \mu_0, \phi - \phi_0) = \left| r_1(-\mu, \mu_0, \phi - \phi_0) - r_{3,4}(-\mu, \mu_0, \phi - \phi_0) \right| \tag{44}$$

and which depends on surface albedo, atmospheric constituent profile, etc., is not used in the cloud detection criterion. Rather, the cloud detection criterion is based on the BDAS angular shape alone. The BDAS algorithm states that high cloud is present if, within measurement uncertainty, the change in BDAS with scattering angle exceeds a threshold value, when calculated using the BDAS reference and comparison cameras. The choice of which cameras to use in this calculation is made by looking at the availability of BRF and angular data, as well as which camera(s) view forward scattered light (determined by calculating the scattering angle, Ω, and choosing the direction for which $\Omega < \Omega'$, where Ω' is a value slightly larger than 90°. The reason for choosing a value slightly larger than 90° is to insure that at least one BDAS retrieval is available near the solar equator; otherwise, the ASCM will have a gap in the vicinity of the solar equator when the retrievals are projected to the SOM grid. If the scattering angle corresponding to both reference cameras satisfy this criterion, then both values of ASCM are determined and projected to the SOM grid).

When choosing a reference camera we first check to see if at least one of the D cameras has valid BRF data in the required bands and is viewing forward scattered light. If this condition is true, we set the value of the reference camera(s) accordingly and stop the search for a reference camera. If neither D camera satisfies the above condition we then check both of the C cameras. If the search still fails, a BDAS retrieval is not possible for this pixel. If one or both of the D cameras is chosen as the reference camera, we then have to search for a comparison camera that has valid BRF and angular data. If Df is chosen as the reference camera, we first check the Cf and then the Bf cameras in our search for a comparison camera. If Cf satisfies the condition for a comparison camera, Bf is not checked. If Cf is chosen as the reference camera, we only check Bf. Similar logic applies to Da, Ca, and Ba.

More formally, defining Ψ:

$$\Psi(\Omega) = \frac{\partial \Delta r}{\partial \Omega} \tag{45}$$

then high cloud is deemed detected if $\Psi \geq \Psi_{thresh}$, where Ω is defined by

$$\cos\Omega = -\mu\mu_0 + (1-\mu^2)^{\frac{1}{2}}(1-\mu_0{}^2)^{\frac{1}{2}}\cos(\phi-\phi_0) \quad (46)$$

Inputs from the Angular Signature Threshold Dataset are employed to choose the appropriate value of Ψ_{thresh}. Nominally, if the cameras are viewing snow/ice-free land, the band 1 - 3 difference is used; if water or snow/ice, the band 1 - 4 difference is used. (The usage of the band 1 - 3 difference or the band 1 - 4 difference is designed to be a configurable input based on surface type, so that the choice may be changed after launch, if desired.) The result of the BDAS test then returns No Retrieval, CloudHC, CloudLC, ClearLC, or ClearHC for each 1.1-km sample. As described in a later section, these are projected to the SOM grid to establish the ASCM. A designation of No Retrieval occurs when any of the four BRF values required for calculation of Eq. (45) is not available.

3.5.2 Generate histograms

3.5.2.1 Physics of the problem

Over land, the thresholds in the Angular Signature Threshold Dataset are updated on a seasonal basis. The updated thresholds are derived from histograms of the BDAS observable, generated from data gathered over the course of a season. An Angular Signature Histogram Dataset is generated for each season. There exist separate Threshold and Histogram datasets for each combination of reference-comparison camera (D-C, D-B and C-B). For each surface type, the BDAS observable, ζ, as defined by Eq. (45), is calculated and stored in the Histogram Dataset(s) in bins, partitioned as follows:

(1) There are 10 sun angle cosine (μ_0) bins: $0.0 \le \mu_0 < 0.1; 0.1 \le \mu_0 < 0.2; ...; 0.9 \le \mu_0 \le 1.0$.

(2) There are 12 relative azimuth ($\Delta\phi = \phi - \phi_0$, the azimuth difference between the view and illumination angles) bins: $0° \le \Delta\phi < 15°; 15° \le \Delta\phi < 30°; ...; 165° \le \Delta\phi \le 180°$. Symmetry about the principal plane is assumed, such a value of $\Delta\phi > 180°$ is replaced by $360° - \Delta\phi$.

(3) The histogram of each observable is broken up into 128 gray levels.

The upper and lower bounds of the histograms are dependent on the surface type, which is established by referencing the Ancillary Geographic Product (AGP), the Cloud Screening Surface Classification (CSSC) Dataset, and the external snow/ice cover input. If the land-water identifier from the AGP indicates land and the CSSC Dataset contains water for that same pixel, we then perform a nearest-neighbor search in the CSSC to find the closest pixel in that dataset that contains a land designation and use its identifier. If the search fails to turn up a land classification within a specified search radius, we use a default land class instead.

3.5.3 Update thresholds

3.5.3.1 Physics of the problem

Because of the variability of the Earth's surface, particularly over land, it is desirable to generate thresholds dynamically for certain parts of the cloud detection algorithm. This is done in automated fashion at the DAAC using the histogramming approach described in §3.5.2, as applied to accumulated MISR observations over a period of time. Analyses of the histograms will be performed to insure that they contain sufficient information to derive updated thresholds.

Thresholds are a function of: View zenith angle, sun angle, relative view-sun azimuth angle, time, and surface classification. In general, the time dependence can be of three kinds: (a) time independent (static), (b) dynamic (seasonal), and (c) dynamic on a less than seasonal time scale. The latter is approximately monthly, and on a sliding window. BDAS uses only static and seasonal thresholds, defined as follows:

(1) Static means that the same threshold is used throughout the mission.

(2) Seasonal means that a separate threshold is defined for Winter, Spring, Summer, and Fall (i.e., Jan-Feb-Mar, Apr-May-Jun, Jul-Aug-Sep, Oct-Nov-Dec). Current data in a particular season uses the threshold derived from the previous year's data in the same season.

For generation of the ASCM, there are three thresholds: T1 divides CloudHC from CloudLC, T2 divides CloudLC from ClearLC, and T3 divides ClearLC from ClearHC. They are either static or dynamic on seasonal time scales, depending on surface type, as shown in Table 7.

Table 7: Time dependencies of thresholds in the Angular Signature Threshold Dataset

Observable	Threshold 1	Threshold 2	Threshold 3
BDAS	Water (shallow or deep) or snow/ice: Static Non-snow/ice covered land: Seasonal	Water (shallow or deep) or snow/ice: Static Non-snow/ice covered land: Seasonal	Water (shallow or deep) or snow/ice: Static Non-snow/ice covered land: Seasonal

3.5.3.2 Mathematical description of the algorithm

Derivation of the three thresholds T1, T2, and T3 from the histograms is done in an automated fashion. T2 is derived using the method of Li and Lee *[5]*. It is expected that over the time scale of a season, the histogram will be bimodal. T1 and T3 are calculated by first finding T1(peak) and T3(peak), the histogram maxima on the respective sides of T2. If multiple bins on either side of T2 have the same frequency, the values of T1(peak) and T3(peak) closest to T2 are selected. We then set

$$T1 = T1(\text{peak}) + a\sigma_1 \tag{46a}$$

$$T3 = T3(\text{peak}) + b\sigma_3 \tag{46b}$$

where σ_1 is the standard deviation of the data on the cloudy side of T2 and σ_3 is the standard deviation of the data on the clear side of T2, and a and b are scene-dependent parameters with $a \leq 0$ and $b \geq 0$. A mathematical description of the histogram analysis method is provided in the discussion for the RCCM, which is found in *[M-5]*.

3.5.4 Project to SOM grid

3.5.4.1 Physics of the problem

Determination of which 1.1-km SOM grid point to assign the ASCM proceeds by assuming the cloudy values of the ASCM are at a climatological altitude of high clouds found within the TASC Dataset. The possibility that the height could be different, and therefore that the ASCM will be assigned to the wrong 1.1-km subregion is deemed acceptable for the at-launch product.

The BDAS is calculated from terrain-projected BRF's. With the BDAS comparison and reference cameras referenced to the same point on the terrain, the C and D camera view vectors do not cross the same point at the cloud altitude as depicted in Figure 19. Nonetheless, by applying the BDAS test to the terrain referenced imagery, Di Girolamo *[4]* has shown that the resulting BDAS cloud mask test is assigned to the reference camera. The mask is then projected along the reference view vector from the terrain to the ellipsoid at the climatological high cloud altitude, and then assigned a location on the SOM grid. The same iterative reprojection technique that we used when projecting the terrain-referenced RCCM to the SOM grid is applied here.

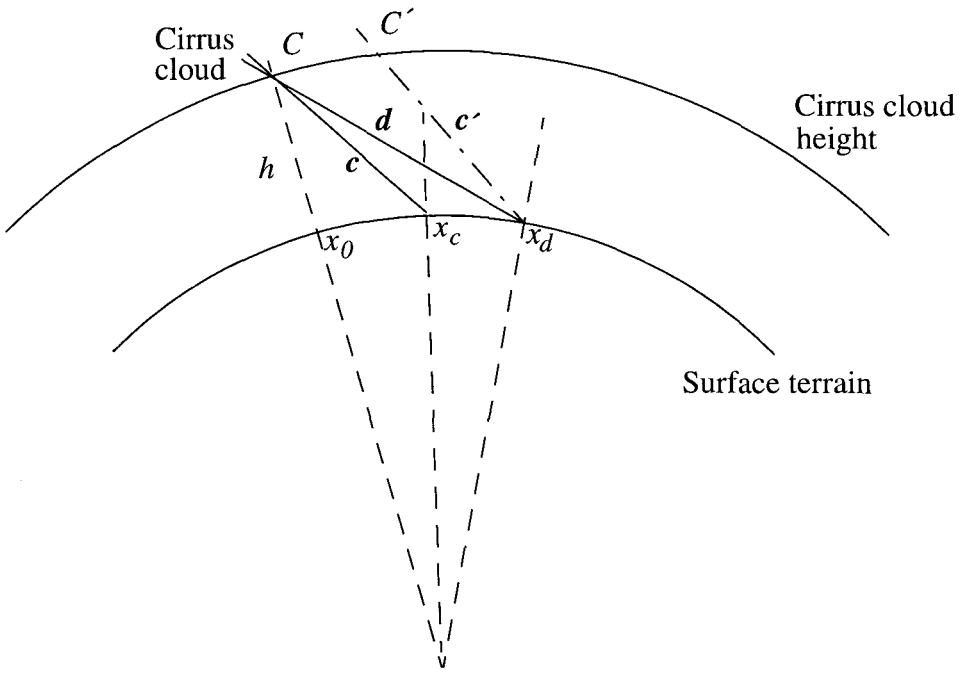

Figure 28. Projection of cirrus cloud from the D and C cameras

3.5.4.2 Mathematical description of the algorithm

We want to project subregions identified as cirrus cloud and clear from the surface terrain onto an ellipsoid at a nominal altitude of h_{cirr}, the climatological high cloud height above the reference WGS84 ellipsoid. This is done using the methodology described in §3.4.7.2, projecting from the terrain to an ellipsoid at the higher elevation $h = h_{cirr}$. In this case the look vector is chosen to be that of the D camera. We then convert the Earth Centered Reference coordinates of this point to an SOM grid location. The mapping from variable terrain altitudes to the high cloud altitudes could result in the following situations:

(1) A single value of the ASCM is contained within a 1.1-km subregion on the ellipsoid. The projected ASCM is assigned this value. The values of the reference and comparison camera used to calculate this mask value are also retained and reprojected to the ellipsoid.

(2) More than one value of the ASCM is contained within the 1.1-km subregion. In this case we assign numerical values of 0, 1, 2, 3 and 4 respectively to the conditions No Retrieval, CloudHC, CloudLC, ClearLC, and ClearHC. The minimum numerical value is then used to select the projected ASCM designation. This condition can occur as a result of terrain effects, or when two subregions, one north and one south of the solar equator, both satisfy the D-camera scattering angle criterion of §3.5.1.2 (nominally, one for Df and the other for Da). The corresponding values of the reference and comparison camera tag along with the prevailing ASCM value.

(3) No value of the ASCM is contained within the 1.1-km subregion. In this case the projected ASCM is designated No Retrieval.

3.5.5 Establish regional scene classification

3.5.5.1 Physics of the problem

The three cloud masks, namely the RCCM, ASCM and the SDCM, and the stereo heights, $\tilde{h}_{ster}$, are used to establish regional scene classifiers. With respect to the RCCM, both the original terrain-projected RCCM and the reference camera RCCM mapped onto the SOM grid as a result of assigning a stereoscopic height (i.e., the feature-projected RCCM) are used. The latter has a snow/ice identifier associated with it. These classifiers contain all the information on the state of cloudiness derived from MISR for regions measuring $(17.6\ \mathrm{km})^2$. They will be useful when intercomparing the relative performance between the different MISR cloud masks and when cross-validating with the cloud masks from MODIS. Also, when coupled with the texture measurements, the regional scene classifiers will provide an excellent framework for higher-level scene classification that can be tuned to meet the needs of different users. The regional scene classifiers consist of the following parameters:

(1) $f^{k}_{m,\,l}$, an altitude-binned classifier in which k represents one of five altitude classes: surface, low cloud, mid-level cloud, high cloud, and no retrieval; m represents one of the three masks (feature-projected RCCM, ASCM, SDCM); and l represents the clear/cloud or surface/cloud and confidence designations associated with each mask. The altitude binning is referenced to the surface terrain.

(2) F_{CloudHC}, the fractional area of the region that is classified, with high confidence, as consisting of any type of cloud, and F_{ClearHC}, the fractional area of the region that is classified, with high confidence, as being clear of any type of cloud.

(3) $\tilde{F}_{\mathrm{CloudHC}}$ and $\tilde{F}_{\mathrm{CloudLC}}$, the fractional area of the region that is classified, with high and low confidence, respectively, as consisting of clouds excluding those detected using the BDAS test. These parameters are used as input to Level 2 aerosol retrievals.

(4) $\hat{F}^k_{\text{CloudHC}}$ and $\hat{F}^k_{\text{CloudLC}}$, the fractional area of the region that is classified, with high and low confidence, respectively, as consisting of clouds detected using the original terrain-projected RCCM only, for camera *k*.

3.5.5.2 Mathematical description of the algorithm

3.5.5.2.1 Regional altitude-binned scene classifiers

Let *R(i,j), A(i,j)* and *S(i,j)* be the output for the $(i,j)^{\text{th}}$ 1.1-km sample of the feature-projected RCCM, ASCM and SDCM, respectively. The designations associated with *R(i,j), A(i,j)* and *S(i,j)* are assigned numerical values, *l*, as shown in Table 8.

Table 8: Numerical values associated with the feature-projected cloud masks

l or *m'*	RCCM	ASCM	SDCM
0	No Retrieval	No Retrieval	No Retrieval
1	CloudHC	CloudHC	CloudHC
2	CloudLC	CloudLC	CloudLC
3	ClearLC	ClearLC	NearSurface
4	ClearHC	ClearHC	Clear

The regional altitude-binned scene classifier, $f^k_{m,l}$, is calculated from:

$$f^k_{m,l} = \frac{1}{256}\sum_{i,j}\delta_{l,m'(i,j)} \cdot \xi^k_{i,j} \tag{47}$$

where δ is the Kronecker delta, which is equal to 1 if $l = m'$ and 0 otherwise. The index *m'(i,j)* is the numerical value associated with the mask *m* in the $(i,j)^{\text{th}}$ sample, and takes on the same set of values shown in Table 8. The variable *k* is an index for the altitude class calculated from the stereoscopic heights. Now, define:

$$\begin{aligned} &\xi^{sfc}_{i,j} = 1 \text{ for } S(i,j) = \text{Clear or NearSurface;} \\ &\qquad\qquad \tilde{h}_{ster}(i,j) \text{ source flag is not Surface;} \\ &\text{otherwise } \xi^{low}_{i,j} = 0 \end{aligned} \tag{48}$$

$\xi_{i,j}^{low} = 1$ for $h_s(i,j) < \tilde{h}_{ster}(i,j) \leq h_{surf}(i,j) + 2$ km;

$S(i,j)$ = CloudLC or CloudHC; and

$\tilde{h}_{ster}(i,j)$ source flag is not Default Cloud or MODIS;

otherwise $\xi_{i,j}^{low} = 0$ (48a)

$\xi_{i,j}^{mid} = 1$ for 2 km $< \tilde{h}_{ster}(i,j) \leq h_{surf}(i,j) + 6$ km;

$S(i,j)$ = CloudLC or CloudHC; and

$\tilde{h}_{ster}(i,j)$ source flag is not Default Cloud or MODIS;

otherwise $\xi_{i,j}^{mid} = 0$ (48b)

$\xi_{i,j}^{high} = 1$ for $\tilde{h}_{ster}(i,j) > h_{surf}(i,j) + 6$ km;

$S(i,j)$ = CloudLC or CloudHC; and

$\tilde{h}_{ster}(i,j)$ source flag is not Default Cloud or MODIS;

otherwise $\xi_{i,j}^{high} = 0$ (48c)

$\xi_{i,j}^{NR} = 1$ for $S(i,j)$ = No Retrieval; or

$S(i,j)$ = CloudLC and $\tilde{h}_{ster}(i,j)$ source flag is Surface, Default Cloud or MODIS;

otherwise $\xi_{i,j}^{high} = 0$ (48d)

where $h_s(i,j)$ is the near-surface threshold defined in §3.4.6.2 and $h_{surf}(i,j)$ is the surface elevation in the 1.1-km sample. Note that the above logic excludes (i.e., treats as No Retrieval) stereoscopic heights that have been modified by overriding the stereoscopic values with Surface, Default Cloud, or MODIS cloud altitudes. This is to insure that the altitude binning of clouds relies solely on purely stereoscopically retrieved values.

3.5.5.2.2 Regional scene classifiers

The total fractional area of the 17.6-km region classified as containing any type of cloud with high confidence is:

$$F_{\text{CloudHC}} = \frac{1}{256}\sum_{i,j}\psi_{i,j}^{CloudHC}[R(i,j), A(i,j), S(i,j)] \quad (49)$$

where the RCCM used is the feature-projected version with its associated snow/ice indicator and:

$\psi^{CloudHC}$ = 1 if $R(i,j)$ or $A(i,j)$ = CloudHC and the surface is free of snow/ice;

$\psi^{CloudHC}$ = 1 if $A(i,j)$ or $S(i,j)$ = CloudHC and the surface is snow/ice covered;

$\psi^{CloudHC}$ = 0 otherwise.

The total fractional area of the 17.6-km region classified as clear of any type of cloud with high confidence is:

$$F_{\text{ClearHC}} = \frac{1}{256}\sum_{i,j}\psi_{i,j}^{ClearHC}[R(i,j), A(i,j), S(i,j)] \tag{50}$$

where the RCCM used is the feature-projected version with its associated snow/ice indicator and:

$\psi^{ClearHC}$ = 1 if $R(i,j) = A(i,j)$ = ClearHC and the surface is free of snow/ice;

$\psi^{ClearHC}$ = 1 if $A(i,j)$ = ClearHC and $S(i,j)$ = Clear and the surface is snow/ice covered;

$\psi^{ClearHC}$ = 0 otherwise.

The estimated fractional area of the 17.6-km region that contains cloud with high confidence, excluding the clouds detected with the ASCM (i.e., with the BDAS test), is:

$$\tilde{F}_{\text{CloudHC}} = \frac{1}{256}\sum_{i,j}\chi_{i,j}^{CloudHC}[R(i,j), S(i,j)] \tag{51}$$

where the RCCM used is the feature-projected version with its associated snow/ice indicator and:

$\chi^{CloudHC}$ = 1 if $R(i,j)$ = CloudHC and the surface is free of snow/ice;

$\chi^{CloudHC}$ = 1 if $S(i,j)$ = CloudHC and the surface is snow/ice covered;

$\chi^{CloudHC}$ = 0 otherwise.

The estimated fractional area of the 17.6-km region that contains cloud with low confidence, excluding the clouds detected with the ASCM, is:

$$\tilde{F}_{\text{CloudLC}} = \frac{1}{256}\sum_{i,j}\chi_{i,j}^{CloudLC}[R(i,j), S(i,j)] \tag{52}$$

where the RCCM used is the feature-projected version with its associated snow/ice indicator and:

$\chi^{CloudLC} = 1$ if $R(i,j)$ = CloudLC and the surface is free of snow/ice;

$\chi^{CloudLC} = 1$ if $S(i,j)$ = CloudLC and the surface is snow/ice covered;

$\chi^{CloudLC} = 0$ otherwise.

The parameters $\tilde{F}_{\text{CloudHC}}$ and $\tilde{F}_{\text{CloudLC}}$ are used in Level 2 aerosol processing to establish whether a region is too cloud-covered to perform a retrieval of the aerosol parameters.

3.5.5.2.3 Regional angle-by-angle cloud fractions

Regional angle-by-angle cloud fractions are calculated using only the terrain-projected RCCM as input. Let $R^k(i,j)$ be the output for the $(i,j)^{\text{th}}$ 1.1-km sample of this RCCM for camera k. Then, the total fractional area of the 17.6-km region classified as cloud with high confidence for the angle-by-angle classifier is defined as

$$\hat{F}^k_{\text{CloudHC}} = \frac{1}{256}\sum_{i,j}\eta^{CloudHC}_{i,j}[R^k(i,j)] \tag{53}$$

where:

$\eta^{CloudHC} = 1$ if $R^k(i,j)$ = CloudHC;

$\eta^{CloudHC} = 0$ otherwise.

Similarly, the total fractional area of the 17.6-km region classified to be cloud with low confidence for the angle-by-angle classifier is defined as

$$\hat{F}^k_{\text{CloudLC}} = \frac{1}{256}\sum_{i,j}\eta^{CloudLC}_{i,j}[R^k(i,j)] \tag{54}$$

where:

$\eta^{CloudLC} = 1$ if $R^k(i,j)$ = CloudLC;

$\eta^{CloudLC} = 0$ otherwise.

3.5.6 Generate Cloud Shadow Mask and Topographic Shadow Mask

3.5.6.1 Physics of the problem

The stereoscopic heights, the SDCM, the SOM-projected version of the RCCM with its associated snow/ice indicator, and the AGP are used to derive a Cloud Shadow Mask (CSM) and a Topographic Shadow Mask (TSM). These will be used as input to Level 2 Aerosol/Surface processing because the presence of shadow violates the assumptions underlying the radiative transfer model applied to the aerosol retrievals. The clouds and topography are projected onto the surface terrain along the direction of the solar illumination vector. We establish a value for the vertical thickness of the clouds. The 1.1-km SOM grid points nearest to locations along the projected ray are flagged as being cloud shadowed if the local terrain altitude lies between the projected altitudes of the cloud bases and tops, and topographically shadowed if the local terrain altitude lies below the projected altitude of the ray-cast topography.

Since we do not incorporate an independent source of data on cloud thickness, a default value of 275 m is chosen. This choice of a low value means that certain subregions that should have been classified as shadowed will be designated as shadow-free. On the other hand, a large value of cloud thickness may be too conservative. At present, we adopt the value indicated on the premise that the cloud shadow mask is a supplement to other screens employed during aerosol retrievals, such as cloud masks and thresholds on the acceptability of aerosol models to fit the observations. The data processing software is designed to enable this value to be changed if necessary.

The CSM and TSM are 1.1-km resolution SOM fields. The CSM has the following designations at each point: Cloud Shadowed with High Confidence (CloudShadowHC), Cloud Shadowed with Low Confidence (CloudShadowLC), or Cloud Shadow Free. The TSM has the following designations at each point: Topographically Shadowed or Topographic Shadow Free. Initially, we consider all locations in the CSM to be Cloud Shadow Free and all locations in the TSM to be Topographic Shadow Free. Then, redesignations as cloud shadowed or topographically shadowed are established according to the procedure presented below. Note that CSM and TSM are independent masks, that is, the same SOM location may be designated both cloud shadowed and topographically shadowed, for example.

3.5.6.2 Mathematical description of the algorithm

We consider for projection only those 1.1-km SOM grid locations that are not classified as water in the AGP (for the purpose of establishing topographic shadows), as well as all land and water locations where $\chi^{CloudHC} = 1$ or $\chi^{CloudLC} = 1$, where $\chi^{CloudHC}$ and $\chi^{CloudLC}$ are defined in Eqs. (51) and (52) and the text following these equations (for the purpose of establishing cloud shadows). The solar illumination vector at a particular location is defined by a unit vector $\hat{s}$, which points in the direction of photon travel. The direction of this vector changes much more slowly over

horizontal scales than camera look vectors, and is thus defined on a coarse grid, e.g., every 17.6 km. For a given 1.1-km grid location, it is sufficient to use the nearest defined solar vector.

For each cloud to be projected onto the surface terrain, we define a cloud top altitude, h_t, and take it to be equal to the stereoscopically derived height, that is, $h_t = \tilde{h}_{ster}$. We also define a nominal cloud thickness, Δh_1, defined by

$$\Delta h_1 = min(\Delta h_{nom}, \Delta h_2) \tag{55}$$

where Δh_{nom} is a nominal cloud thickness (presently set to 275 m) and $\Delta h_2 = |h_t - h_{surf}|$, where h_{surf} is the surface terrain elevation at the SOM grid point where the cloud is located. In other words, we do not permit the cloud to extend below the surface.

At each point where projection is to be performed, we ray cast from an altitude h_0, defined to be the cloud top altitude, h_t (if cloudy) or the topographic elevation, h_{surf} (if clear). We now need to march along the solar vector projected from this altitude to find out which points are in shadow. Since the solar vector is given in ECR coordinates, we need to establish a transformation to SOM grid space. To do this, we use the methodology described in §3.4.7.2, but replace the camera look vector $\hat{l}$ by the solar vector $\hat{s}$, and in Figures 23 and 24 interpret location B to be at h_0 and location E to be a point on the surface ellipsoid. In solving for the distance between these two points, λ, we choose the solution represented by Eq. (38), with the normalized solar vector $\hat{s}'$ replacing $\hat{l}'$. Then

$$\hat{e} = \hat{p} + \lambda \hat{s} \tag{56}$$

describes the location where the solar illumination vector intersects the surface ellipsoid. We translate this point from the geocentric coordinates to SOM coordinates. The projection of the solar vector on the SOM x-y plane can then be defined by a unit vector as:

$$\hat{s}_{som} = \begin{bmatrix} \dfrac{x_E - x_B}{\lambda_{som}} \\ \dfrac{y_E - y_B}{\lambda_{som}} \end{bmatrix} \tag{57}$$

where (x_E, y_E) and (x_B, y_B) are the SOM coordinates of points E and B, and λ_{som} is defined by:

$$\lambda_{som} = \sqrt{(x_E - x_B)^2 + (y_E - y_B)^2} \tag{58}$$

Assuming each step in the SOM x-y plane is $\Delta\lambda_{som} = 550$ m for the ray casting, then the maximum

number of steps required to step from the cloud top to its intersection with the surface ellipsoid is given by:

$$N_{steps} = \frac{\lambda_{max}}{550} \tag{59}$$

with λ_{max} given in meters and defined by the smaller of λ_{som} and a maximum shadow-casting distance of 35.2 km. The distance of each step along the solar ray vector in the three-dimensional geocentric coordinate system is determined by the total distance between the locations E and B, and the total number of steps:

$$\Delta\lambda_{ray} = \frac{\lambda}{N_{steps}} \tag{60}$$

The ray-casting direction is constant over $(17.6\text{-km})^2$ regions. Once the ray-casting direction on the SOM grid has been determined, we start from the coordinates $(x_0, y_0)_{som}$ representing the center of each 1.1-km SOM subregion to be projected. Let the line and sample coordinates of these starting locations be written as $(l, s)_0$. We will do the ray casting from $(x_0, y_0)_{som}$ along the direction defined by the unit vector $\hat{s}_{som}$ up to a maximum stepping distance λ_{max}. At each step point $(x_i, y_i)_{som}$, determined by stepping 550 m along the direction defined by $\hat{s}_{som}$, the following procedure is followed:

(1) Find the nearest line and sample coordinates $(l, s)_i$ to the point $(x_i, y_i)_{som}$ and find the terrain elevation at this SOM grid location, $h_{terrain,i}$;

(2) Compute the current solar ray elevation at the step point. The three-dimensional geocentric coordinate of the point is derived by $B_i = B + i\Delta\lambda_{ray}\hat{s}$, where i is the number of steps. Converting it into geodetic coordinates will result in the current altitude h_i of the solar ray passing through point B;

(3) If $\chi^{CloudHC}$ or $\chi^{CloudLC}$ of the projected point = 1, set the ray altitude corresponding to the cloud top according to $h_{t,i} = h_i$, and establish the ray altitude corresponding to the cloud base by $h_{b,i} = h_{t,i} - \Delta h_1$. Set the ray altitude corresponding to the projected topography to $h_{topog,i} = h_{t,i} - \Delta h_2$. Otherwise, set $h_{topog,i} = h_i$;

(4) If $\chi^{CloudHC}$ or $\chi^{CloudLC}$ of the projected point = 1, and $h_{b,i} \le h_{terrain,i} \le h_{t,i}$, set the CSM of $(l, s)_i$ to:

CloudShadowHC if $\chi^{CloudHC} = 1$;

CloudShadowLC if $\chi^{CloudLC} = 1$, but only if the CSM of $(l, s)_i$ is not already designated CloudShadowHC, in which case the CloudShadowHC designation is retained;

(5) If $h_{terrain,i} \leq h_{topog,i}$, set the TSM to Topographically Shadowed;

(6) Continue stepping until $h_i < h_{terrain,\,i}$ or until $i = N_{steps}$. Repeat the above evaluations at each new step.

3.6 PRACTICAL CONSIDERATIONS

3.6.1 Numerical computation considerations

Requirements on processing speed and data storage are described in *[M-17]*.

3.6.2 Programming and procedural considerations

Guidelines to be followed during algorithm development are described in *[M-13]*.

3.6.3 Configuration of retrievals

A Cloud Classification Configuration File is used to establish the numerical values of adjustable parameters used within the retrievals, e.g., thresholds establishing whether a successful stereo matching occurred. The purpose of establishing a separate file is to avoid "hard-wiring" specific values into the software. The TOA/Cloud Product will contain information indicating what version of the configuration file was used. The contents of the Cloud Classification Configuration File are shown in Table 9. The values shown correspond to the at-launch settings. The column entitled "Section" indicates where in this ATB a description of the specific configuration parameter is found.

Table 9: Contents of the Cloud Classification Configuration File

Parameter	Value	Section
Maximum acceptable RDQI used in averaging data to appropriate resolution, $RDQI_1$	1	3.3.1.2
$RDQI_2$ [value of RDQI' in Eq. (0b) if RDQI > $RDQI_1$]	3	3.3.1.2
Minimum solar zenith angle cosine for conversion of radiances to BRF's	0.01	3.3.2.2
Nominal along-track view zenith angle for camera Df	70.5°	3.4.1.1
Nominal along-track view zenith angle for camera Cf	60.0°	3.4.1.1
Nominal along-track view zenith angle for camera Bf	45.6°	3.4.1.1
Nominal along-track view zenith angle for camera Af	26.1°	3.4.1.1
Nominal along-track view zenith angle for camera An	0.0°	3.4.1.1
Nominal along-track view zenith angle for camera Aa	-26.1°	3.4.1.1
Nominal along-track view zenith angle for camera Ba	-45.6°	3.4.1.1

Table 9: Contents of the Cloud Classification Configuration File (continued)

Parameter	Value	Section
Nominal along-track view zenith angle for camera Ca	-60.0°	3.4.1.1
Nominal along-track view zenith angle for camera Da	-70.5°	3.4.1.1
Nominal along-track time difference between cameras Cf and Df (lines)	1468	3.4.1.1
Nominal along-track time difference between cameras Bf and Cf (lines)	1292	3.4.1.1
Nominal along-track time difference between cameras Af and Bf (lines)	1127	3.4.1.1
Nominal along-track time difference between cameras An and Af (lines)	1113	3.4.1.1
Nominal along-track time difference between cameras Aa and An (lines)	1113	3.4.1.1
Nominal along-track time difference between cameras Ba and Aa (lines)	1127	3.4.1.1
Nominal along-track time difference between cameras Ca and Ba (lines)	1292	3.4.1.1
Nominal along-track time difference between cameras Da and Ca (lines)	1468	3.4.1.1
Number of stereo matcher pyramid levels	2	3.4.1.2
RDQI threshold for determining what constitutes valid data for stereo retrievals, $RDQI_3$	1	3.4.1.2.1, 3.4.1.2.2, 3.4.1.3.1, 3.4.1.3.2
~~Minimum percentage of points needed for match in M2, M3~~	~~75%~~	
Factor for setting secondary threshold in ambiguity test, f	1.1	3.4.1.2.1, 3.4.1.2.2
Along-track spatial clustering threshold, p_x	3 pixels	3.4.1.2.1, 3.4.1.2.2
Cross-track spatial clustering threshold, p_y	3 pixels	3.4.1.2.1, 3.4.1.2.2
Fraction of pixels required to be valid ($RDQI \leq RDQI_3$) in order to use NM, f_{NM}	0	3.4.1.3.1
Highest usable level in NM	5th	3.4.1.3.1
Lowest usable level in NM	2nd lowest	3.4.1.3.1
Search radius used in NM, when not set to 1	2	3.4.1.3.1
Highest level at which search radius other than 1 is used	1	3.4.1.3.1
Maximum cloud height for establishing search window and working strip dimensions in stereo matchers	20 km	3.4.1.3.1, 3.4.3.2, 3.4.5.2
Maximum cloud velocity for establishing search window and working strip dimensions in stereo matchers	100 m/sec	3.4.1.3.1, 3.4.3.2, 3.4.5.2

Table 9: Contents of the Cloud Classification Configuration File (continued)

Parameter	Value	Section
Search window pad	1.1 km	3.4.1.3.1
Threshold q_{thresh} in stereo matcher applicability test	5	3.4.1.4
~~Minimum amount of BRF data needed for each camera in Stage 2~~	~~80%~~	
~~Maximum number of No Retrievals allowable in RCCM in Stage 2~~	~~20%~~	
~~Proximity requirement for height and wind retrieval camera pairs~~	~~≤ 3~~	
~~Wind retrieval camera determinant threshold~~	~~200~~	
Height retrieval forward camera pair	An, Af	3.4.2
Height retrieval aftward camera pair	An, Aa	3.4.2
Wind retrieval forward camera triplet	An, Bf, Df	3.4.2
Wind retrieval aftward camera triplet	An, Ba, Da	3.4.2
M2 patch size used in wind retrievals, M_c x M_a	11 pixels x 21 pixels	3.4.3.2
Minimum allowable height for use in wind retrieval	0 km	3.4.4.2
Maximum allowable height for use in wind retrieval	20 km	3.4.4.2
Width of wind speed histogram bins	6 m/sec	3.4.4.2
Minimum value of wind speed in histogram, along-track and cross-track	-102 m/sec	3.4.4.2
Maximum value of wind speed in histogram, along-track and cross-track	102 m/sec	3.4.4.2
Minimum number of wind retrieval points required to avoid defaults	100	3.4.4.2
Default value of h_{wind} if wind retrieval fails	10 km	3.4.4.2
Default along-track and cross-track wind speeds if wind retrieval fails	0 m/sec	3.4.4.2
Wind speed differential for merging of search windows	12 m/sec	3.4.5.2
Fraction of data in search window required to be available, f_{search}	0.75	3.4.5.2
Search window pad dimension, l	1.1 km	3.4.5.2.1
Target patch size used in height retrievals, N_c x N_a	6 pixels x 10 pixels	3.4.5.2.1
Small search window dimensions when using previous disparities, B_c x B_a	5 pixels x 7 pixels	3.4.5.2.1
M2 matcher factor f_2	0.5	3.4.5.2.1
M3 matcher factor f_3	0.5	3.4.5.2.1
M2 threshold	0.75	3.4.5.2.2
M3 threshold	1.0	3.4.5.2.2
Search window size at second level of pyramid, $L_c' x L_a'$	5 pixels x 5 pixels	3.4.5.2.2
RS filter window size, J_c x J_a	6 pixels x 6 pixels	3.4.5.2.4

Table 9: Contents of the Cloud Classification Configuration File (continued)

Parameter	Value	Section
RS along-track string size, K_a	7 pixels	3.4.5.2.4
Required minimum population of filter window to perform median filtering	50%	3.4.5.2.4
Matcher score, M2 verified by M3	4	3.4.6.2
Matcher score, M2	3	3.4.6.2
Matcher score, M3	2	3.4.6.2
Matcher score, RS	1	3.4.6.2
Minimum allowable stereoscopic height	0 km	3.4.6.2
Maximum allowable stereoscopic height	20 km	3.4.6.2
Ray Skewness Test Blunder threshold	1500 m	3.4.6.2
Ray Skewness Test Blunder score	0	3.4.6.2
Ray Skewness Test Low Confidence score	1	3.4.6.2
Ray Skewness Test High Confidence threshold	275 m	3.4.6.2
Ray Skewness Test High Confidence score	2	3.4.6.2
Stereoscopic height uncertainty increment h_{unc}	275 m	3.4.6.2
Numerical score for declaring a stereoscopic height High Confidence	6 - 8	3.4.6.2
Numerical score for declaring a stereoscopic height Low Confidence	1 - 5	3.4.6.2
Blunder test threshold factor N_σ	2.0	3.4.6.2
Equivalent water threshold for snow cover	5 mm	3.4.7.1
Threshold for sea ice cover areal fraction	5%	3.4.7.1
Default cloud height	3 km above terrain	3.4.8.2
Number of 1.1-km samples in terrain "neighborhood"	7 x 7	3.4.8.2
Decision matrix for establishing final stereoscopic heights	Table 5	3.4.8.2
Classifications associated with the final stereoscopic heights	Table 6	3.4.8.2
RDQI threshold for determining what constitutes valid data in calculating average BRF's and texture indices, $RDQI_4$	0	3.4.11.1
Scattering angle threshold for establishing forward scattering, Ω'	92°	3.5.1.2
Bands used in BDAS test	Based on surface type	3.5.1.2
High cloud-top altitude	TASC climatology	3.5.4.1
Low cloud upper limit for altitude binning	2 km	3.5.5.2
Middle cloud upper limit for altitude binning	6 km	3.5.5.2

Table 9: Contents of the Cloud Classification Configuration File (continued)

Parameter	Value	Section
Cloud thickness used in cloud shadow algorithm	275 m	3.5.6.2
Horizontal ray-casting step size in cloud and topographic shadow algorithm	550 m	3.5.6.2
Maximum horizontal distance for casting cloud and topographic shadow	35.2 km	3.5.6.2

3.6.4 Quality assessment and diagnostics

Several parameters will be reported as part of the TOA/Cloud Product which will serve as diagnostics of how the cloud classification is performing. Maps or other summaries of these parameters will be reviewed by the MISR team for quality assessment purposes. These parameters are described in *[M-18]*. References to sections in this ATB containing descriptions of specific retrieval quality indicators are provided in that document.

3.6.5 Exception handling

It is possible that data in one or more of the 36 instrument channels could be missing, e.g., due to high values of the RDQI or to partial or complete failure of a camera in flight. These cases have been handled throughout this document either through the provision of default parameters or through retrieval approaches using alternative channels or cameras.

3.7 ALGORITHM VALIDATION

Details on planned field campaigns, experimental methodologies, and instrument calibration and data reduction procedures are documented in *[M-16]*.

3.8 ALGORITHM DEVELOPMENT SCHEDULE

A strategy for time-phased development of the algorithms for the products and datasets described in this documents, and a listing of key development milestones, are provided in *[M-13]*.

4. ASSUMPTIONS AND LIMITATIONS

4.1 ASSUMPTIONS

The following assumptions are made with respect to cloud classifications described in this document:

(1) MODIS Level 2 cloud masks and cloud-top heights will be available for incorporation into the systematic processing of MISR data.

(2) Snow-cover and ice-cover products based on passive microwave data will be available from NSIDC.

(3) Vertical cloud motions are ignored for the purposes of making wind displacement corrections in the RLRA retrievals.

(4) Horizontal wind is assumed constant for a given altitude over distances of ~100 km.

4.2 LIMITATIONS

The following limitations apply to the at-launch cloud classifications described in this document:

(1) The reported RLRA is not strictly synonymous with cloud-top height. It is the altitude assigned to the height at which the maximum horizontal contrast in the reflectivities originates. This may be the terrain surface or the tops of bright clouds, but thin clouds can still be present above the RLRA.

(2) Cloud fractions reported as middle and high do not preclude the presence of additional lower-level clouds that may not have been observed due to obscuration by the higher cloud.

(3) No attempt is presently included in the processing to propagate side-hitting radiances through potentially clear portions of the vertical columns for the reflecting level parameters. This will cause some data to be lost in an unbiased way for these parameters, with gains in processing efficiency. The side-hitting radiances are nevertheless used in calculating the coarse (35.2 km) resolution parameters described in *[M-9]*.

(4) Larger than usual errors in the RLRA can be expected in the presence of strong horizontal wind shear (e.g., fronts, jet streams, hurricanes) for which assumption (4) above is invalid.

(5) Rapid changes in surface reflectivities, e.g., due to agricultural changes, local flooding, etc. over land will temporarily affect the appropriate thresholds for cloud detection.

(6) Stratospheric volcanic aerosol will appear as high cloud if it is thick enough.

5.REFERENCES

[1] Chen, D.W., S.K. Sengupta, and R.M. Welch (1989). Cloud field classification based upon high spatial resolution textural features. 2. Simplified vector approach. *J. Geophys. Res.* **97**, 14749.

[2] Di Girolamo, L. (1992). On the detection of cirrus clouds from satellite measurements. Master of Science thesis, Dept. of Meteorology, McGill Univ., Montreal, PQ.

[3] Di Girolamo, L. and R. Davies (1994). A Band-Differenced Angular Signature technique for cirrus cloud detection. *IEEE Trans. Geosci. Rem. Sens.* **32**, 890-896.

[4] Di Girolamo, L. (1996). Detecting and interpreting clouds from satellite radiometric measurements with application to the Multi-angle Imaging SpectroRadiometer (MISR). Ph.D. Thesis, Dept. Atmos. Ocean. Sci., McGill Univ., Montreal, PQ.

[5] Li, C.H. and C.K. Lee (1993). Minimum cross entropy thresholding. *Pattern Recogn.* **26**, 617-625.

[6] Schiffer, R. A. and W. B. Rossow (1983). The International Satellite Cloud Climatology Project (ISCCP): The First Project of the World Climate Research Program (WCRP). *Bull. Amer. Met. Soc.* **64**, 779.

[7] U.S. Geological Survey (1987). *Map Projection - A Working Manual*, Professional Paper 1395, pp. 214-229.